Flip-width: Cops and Robber on dense graphs

Szymon Toruńczyk*

February 2, 2023

Abstract

We define new graph parameters that generalize tree-width, degeneracy, and generalized coloring numbers for sparse graphs, and clique-width and twin-width for dense graphs. Those parameters are defined using variants of the Cops and Robber game, in which the robber has speed bounded by a fixed constant $r \in \mathbb{N} \cup \{\infty\}$, and the cops perform flips (or perturbations) of the considered graph. We propose a new notion of tameness of a graph class, called *bounded flip-width*, which is a dense counterpart of classes of bounded expansion of Nešetřil and Ossona de Mendez, and includes classes of bounded twin-width of Bonnet, Kim, Thomassé and Watrigant. We prove that boundedness of flip-width is preserved by first-order interpretations, or transductions, generalizing previous results concerning classes of bounded expansion and bounded twin-width. We provide an algorithm approximating the flip-width of a given graph, which runs in slicewise polynomial time (XP) in the size of the graph. We also propose a more general notion of tameness, called *almost bounded flip-width*, which is a dense counterpart of nowhere dense classes, and includes all structurally nowhere dense classes. We conjecture, and provide evidence, that classes with almost bounded flip-width coincide with monadically dependent classes, introduced by Shelah in model theory.

*University of Warsaw, Poland. This paper is part of projects that have received funding from the European Research Council (ERC) (grant agreement No 948057 – BOBR).

Contents

1 Introduction

A recent focus of algorithmic and structural graph theory, and of finite model theory, is to find graph parameters that extend the parameters used in the context of sparse graphs, to the dense setting. More generally, the goal is to extend the Sparsity theory of Nešetřil and Ossona de Mendez [NdM12] to dense graph classes. Two central parameters used in the sparse setting are treewidth and degeneracy (also called the coloring number); both have found numerous applications in algorithms and combinatorics. Whereas the dense analogue of treewidth – clique-width, or rank-width – is well-understood, there is not even a candidate for the dense analogue of degeneracy. Recall that a graph has degeneracy at most d if it admits a total order on the vertices such that every vertex has at most d smaller neighbors. Generalized coloring numbers are related parameters, specified by a radius r, which impose restrictions on neighborhoods of radius r, degeneracy being the case of radius $r = 1$. Sparsity theory is a very successful theory studying classes of sparse graphs (in particular, having $O(n^{1+\varepsilon})$ edges for any fixed $\varepsilon > 0$, where n is the number of vertices) in which degeneracy and generalized coloring numbers play a central role.

One of the driving questions in this line of work, on the algorithmic side, is to characterize the *hereditary* (closed under induced subgraphs) graph classes for which the model-checking problem for first-order logic is *fixed-parameter tractable*: there is an algorithm that determines whether a given graph G from the class satisfies a given first-order sentence φ in time $f(\varphi) \cdot |G|^c$, for some constant c and function f that depend only on the class. In the special case of *monotone* classes (closed under subgraphs), Sparsity theory delivered the right notion of combinatorial tractability – that of *nowhere dense* graph classes – which was shown to exactly capture the classes with tractable first-order model-checking by Grohe, Kreutzer and Siebertz [GKS17a]. Alongside with nowhere denseness, the more restrictive graph classes with *bounded expansion* play a central role in the theory. They are in some ways better behaved and simpler to understand than nowhere dense classes, while providing key insights to the latter. A class of graphs $\mathcal{C}$ has bounded expansion if and only if each generalized coloring number is bounded by a constant, on all graphs $G \in \mathcal{C}$. Nowhere dense classes are characterized in an analogous way, with the constant bound replaced with $O(n^{\varepsilon})$, for any fixed $\varepsilon > 0$, where n is the number of vertices of the considered graph G. Those two characterizations play a central role in the model-checking problem, and they have numerous other algorithmic and combinatorial applications, see e.g. [BGMR22].

There is an ongoing effort to lift the notions and results of Sparsity theory to hereditary classes. It is conjectured[1] (see [war16, GHO$^+$20]) that first-order model-checking is fixed-parameter tractable on a hereditary graph class $\mathcal{C}$ if and only if $\mathcal{C}$ is *monadically dependent* (also called *monadically NIP*). This is a notion originating in model theory and introduced by Shelah [She86] (see also [BL21]) in his momentous classification program of logical theories. Intuitively, a hereditary class $\mathcal{C}$ is monadically dependent if one cannot encode arbitrary graphs in graphs in $\mathcal{C}$, using any fixed first-order formula $\varphi(x, y)$.

[1]The conjecture has been circulating in the community since around 2016. As far as we know, it has first been stated explicitly during the open problem session of [war16]. There, dependent (or NIP) classes were considered instead of monadically dependent classes, but those two notions coincide for hereditary classes, by a result of Braunfeld and Laskowski [BL22].

Conjecture 1.1 ([war16])**.** *Let $\mathcal{C}$ be a hereditary class of graphs. Then the model-checking problem for first-order logic is fixed-parameter tractable on $\mathcal{C}$ if and only if $\mathcal{C}$ is monadically dependent.*

Monadically dependent graph classes include all nowhere dense classes [AA14, PZ78], as well as many dense graph classes, such as classes of bounded twin-width [BKTW20]. Monadically dependent graph classes are considered (see [BL21, AA14, NdMP+21, GPT22]) to be the 'correct' dense analogue of nowhere dense classes (as expressed e.g. by Conjecture 1.1), as nowhere dense classes can be characterized, in logical terms, as precisely the monadically dependent classes that exclude some biclique as a subgraph. However, a combinatorial understanding of monadically dependent graph classes, extending the theory developed by Nešetřil and Ossona de Mendez for nowhere dense classes, is lacking. While some fragments of the theory have been lifted (see e.g. [NdMP+21, BL22, DMST22]), the combinatorial core of the theory is still missing, as the most fundamental questions remain unanswered: What is the dense analogue of degeneracy? Of generalized coloring numbers? Of classes with bounded expansion? We believe that these questions need to be answered in order to lift the main algorithmic and combinatorial results of Sparsity theory from the sparse (or monotone) setting, to the dense (or hereditary) setting.

It is not clear how to precisely pose those questions, though. The dense counterpart of the notion of bounded expansion should include all classes with bounded expansion, as well as classes of bounded clique-width, as those are the 'correct' dense analogue of classes of bounded tree-width. It should be contained in monadically dependent classes. We should additionally require that the proposed notion has some good closure properties, and some signs of being combinatorially and algorithmically tame. Perhaps due to the open-ended nature of the question, it hasn't been explicitly stated in the literature[2], nevertheless it has been a driving factor, and a major open problem in the area (see Related Work below). The expectation is that the *right* dense analogues of degeneracy, generalized coloring numbers, and of classes with bounded expansion, will be recognized, once a compelling combinatorial notion, together with results relating it to known classes, is provided. To date, no such notions have been proposed.

We remark that the recently introduced, and already very successful graph parameter *twin-width* [BKTW20] does generalize clique-width, in the sense that classes of bounded clique-width have bounded twin-width. Also, many well-studied classes of bounded expansion – such as classes excluding a fixed minor – have bounded twin-width. However, twin-width does not generalize bounded expansion. Indeed, one of the simplest graph classes from the perspective of Sparsity theory – the class of subcubic graphs, or graphs of maximum degree three – has unbounded twin-width, but has bounded expansion. It is expected that the dense analogue of classes of bounded expansion should include classes of bounded twin-width (see e.g. [GPT22, Section 4],[BNdMS22a, Section 8]).

Contribution We propose a new family of graph parameters, called *flip-width* of radius r, for $r \in \mathbb{N} \cup \{\infty\}$, that are based on games that are similar to the Cops

[2] A rather elusive attempt of defining such classes has been made in [GPT22, Section 4] (see also [BNdMS22a, Section 8]), but this is a far cry from a useful characterization, and its aim is mostly conjectural. See also the section on related work below.

and Robber game considered by Seymour and Thomas to characterize treewidth, in their classic paper [ST93]. Variants of our game can be used to characterize all the aforementioned parameters: treewidth, degeneracy, and generalized coloring numbers used in the context of sparse classes, as well as clique-width and twin-width in the context of dense graph classes. More importantly, they provide generalizations of degeneracy and generalized coloring numbers, to the setting of dense graphs, and offer a compelling dense counterpart of classes of bounded expansion: classes of bounded *flip-width*. We conjecture that classes with *almost bounded flip-width* are precisely the monadically dependent classes. This would give a combinatorial, quantitative characterization of monadic dependence, analogous to the characterization of nowhere dense classes in terms of generalized coloring numbers.

Our starting point is the – apparently new – observation that classes with bounded expansion can be characterized in terms of a variant of the Cops and Robber game. Recall that Seymour and Thomas [ST93] considered a variant of the Cops and Robber in which there are k cops and a robber occupying the vertices of a graph. In each round, some of the cops move in helicopters – that is, not necessarily along edges in the graph – whereas the robber may traverse any path in the graph which avoids the vertices occupied by the cops that remain on ground. The minimal number k of cops needed to capture the robber is equal to one plus the tree-width of the graph.

In our variant of the game the robber runs at speed r, for some fixed $r \in \mathbb{N} \cup \{\infty\}$: he may traverse a path of length at most r that does not run through a cop (he may also stay put). We call this game the *cop-width game* with radius r and width k, if there are k cops and the robber can run at speed r. The *radius-r cop-width* of G, denoted $\mathrm{copwidth}_r(G)$, is the least number k such that the cops win the cop-width game with radius r and width k. (See Sec. 3 for details).

Thus, we obtain a family of graph parameters, one for each $r \in \mathbb{N} \cup \{\infty\}$. As we observe, variants of the cop-width game characterize the graph parameters mentioned earlier: treewidth (for $r = \infty$), degeneracy (for $r = 1$), and generalized coloring numbers (for $1 \leqslant r < \infty$). In particular (see Corollary 3.6), a graph class $\mathcal{C}$ has bounded expansion if and only if[3] $\mathrm{copwidth}_r(\mathcal{C}) < \infty$ for every fixed $r \in \mathbb{N}$. Moreover, a hereditary graph class $\mathcal{C}$ is nowhere dense if and only if for every fixed radius $r \in \mathbb{N}$ and $\varepsilon > 0$, and for all n-vertex graphs $G \in \mathcal{C}$, we have that $\mathrm{copwidth}_r(G) = O(n^\varepsilon)$ (see Corollary 3.7). The cop-width variant of the game, and the graph parameters introduced above, are therefore only suited to the study of sparse graphs.

We introduce a variant of the cop-width game that is suitable for dense graphs, dubbed the *flip-width game*. This is similar to the recent development [GHN+12, GMM+23], where in a certain context, it is shown that vertex removals in sparse graphs correspond to *flips* in dense graphs. The paper [GMM+23] also studies a game based on flips, called the *flipper game*, which should not be confused with the flip-width game introduced below, and served as an inspiration to the flip-width game. The relationship between the two games is discussed in the related work section below.

Flipping a pair of sets $X, Y \subseteq V(G)$ in a graph G results in the graph G' obtained

[3]Here and later on, if f if a graph parameter and $\mathcal{C}$ a graph class, then $f(\mathcal{C})$ denotes the supremum of $f(G)$, for all $G \in \mathcal{C}$.

from G by inverting the adjacency of every pair of vertices $x \in X$ and $y \in Y$ (thus, a pair $x \in X$ and $y \in Y$ is adjacent in G if and only if it is non-adjacent in G'). A *k-flip* of a graph G is obtained by partitioning $V(G)$ into k parts, and then performing flips between some pairs X, Y of parts of the partition (possibly with $X = Y$).

The flip-width game with radius r and width k proceeds in a similar fashion to the cop-width game, but the cops now announce a k-flip of the original graph G in each round, which is then put into effect, in the following sense. When the next k-flip of G is announced by the cops, the robber can traverse a path of length at most r in the *previous* k-flip of G. The game is won by the cops if the new position of the robber is isolated in the announced k-flip. The *flip-width* of radius r of a graph G is the smallest number k needed for the cops to win this variant of the game on G. (See Section 4 for details.) As we prove, variants of the flip-width game characterize clique-width (for $r = \infty$) and twin-width (for $r = 1$, and a variant of flip-width tailored to ordered graphs). Hence, variants of our games capture degeneracy, tree-width, bounded expansion, clique-width, and twin-width, all of which are of central importance in structural and algorithmic graph theory.

The introduction of the flip-width parameters is one of the main contributions of this paper. We supply this notion with an array of results connecting them to other well-known graph parameters, such as the ones listed above. In particular, in Theorem 5.1 we prove that in $K_{t,t}$-free graphs, radius-one flip-width is functionally equivalent to degeneracy, while in ordered graphs (see Theorem 6.3), flip-width is functionally equivalent to twin-width.

The second main contribution of the paper is the notion of a graph class of bounded flip-width, and the results describing their relationship to other graph classes. In particular, we show that classes of bounded flip-width subsume classes with bounded expansion, classes with structurally bounded expansion, and classes of bounded twin-width. Theorem 5.2 says that classes of bounded expansion are exactly those classes of bounded flip-width that exclude some biclique as a subgraph. Furthermore, classes of bounded twin-width are exactly the classes of ordered graphs with bounded flip-width, with the order forgotten (see Corollary 6.4).

The third main contribution are results showing that classes of bounded flip-width enjoy good closure properties, in particular, closure under first-order transductions (see Theorem 7.2). For instance, the class of squares of graphs from a class of bounded flip-width, again has bounded flip-width. This generalizes a prior analogous result for classes of bounded twin-width [BKTW20].

Determining the exact flip-width of radius r of a given graph G seems computationally difficult (the naive approach is exponential in the size of G). As our fourth contribution, we show a slicewise polynomial (XP) algorithm that approximates the flip-width of a given graph G, which runs in time polynomial in the of size, when the flip-width is considered fixed.

As our final contribution, we also introduce the notion of a class of almost bounded flip-width, generalizing classes with bounded flip-width. We prove that they contain nowhere dense classes, and more generally, structurally nowhere dense classes, and conjecture that they coincide with monadically dependent classes.

We supplement our results with a discussion, providing evidence for various stated conjectures, and outline a potential approach towards an algorithmic and combinatorial understanding of the classes introduced in this paper.

Related work The flip-width game is inspired by the paper [GMM+23]. There, another game based on performing flips is introduced, and is called *flipper game*. It is used to characterize monadically stable graph classes. The game differs from the flip-width game in three major ways. In the flipper game, the duration of the game is bounded, whereas in the flip-width game, it is unbounded. On the other hand, in the flip-width game the resources of the flipper are limited by bounding the complexity of the allowed flips. Finally, in each round of the flipper game, the fugitive is confined to the radius-r ball around his present position in the current flip of the graph in all future rounds, not just the next round. The flipper game is used to characterize monadically stable classes, which are incomparable with classes of bounded flip-width, and are likely to be contained in classes of almost bounded flip-width (see Section 9).

There are multiple variations of the Cops and Robber game, and some of them are similar to our cop-width game for sparse graphs. In some of those variants, the robber has bounded speed, as in our cop-width game and flip-width game. See [FKL12, FGK+10, AM15]. For a survey, see [FT08]. The cop-width parameters that we study are closely related to the parameters studied in [RT08] and [LPPT20], which are defined in terms of a similar game, as we discuss in Section 3 and in Appendix A.1.

Some attempts at defining graph classes that are dense analogues of classes of bounded expansion were made in [KPS17, GKN+18, NORS21, NdMP+21, GPT22, BNdMS22a]. In [KPS17, NORS21, NdMP+21], the property of having low rank-width covers is proposed as the dense analogue of bounded expansion. However, this notion does not include[4] classes of bounded twin-width, and it is not known whether this notion is closed under transductions. It has been proved recently that those classes are monadically dependent [BNdMS22b]. In [GPT22, Section 4] and [BNdMS22a, Section 8], an attempt at formalizing an abstract notion of a dense analogue of classes with bounded expansion was made, but those papers don't propose any workable combinatorial definition (however, classes admitting low rankwidth covers do not satisfy the requirements proposed there, as they do not include classes of bounded twin-width).

Acknowledgement I am very indebted to Jan Dreier, Jakub Gajarský, Rose McCarty, Michał Pilipczuk, Wojtek Przybyszewski, and Marek Sokołowski, for numerous stimulating discussions on topics closely related to this paper. Special thanks go to Pierre Ohlmann, who has read a first draft of this manuscript and helped improve it.

Organization The organization of this manuscript is as follows.

Section 2 introduces some standard notation, and recalls some basic notions and results, mostly from Sparsity theory.

In Section 3 we introduce the cop-width parameters, and demonstrate they characterize degeneracy, tree-width, and generalized coloring numbers. Using those parameters, we characterize classes of bounded expansion and nowhere dense classes.

In Section 4 we introduce our two main notions: the flip-width parameters, and

[4] An example was provided by Jakub Gajarský

classes of bounded flip-width. We give many examples, and obtain some combinatorial properties. We show that radius-∞ flip-width is equivalent to clique-width.

In Section 5 we study the behavior of the flip-width parameters in the case of $K_{t,t}$-free graphs. We show that radius-one flip-width corresponds to degeneracy, and that for higher radii, the flip-width parameters correspond to the generalized coloring numbers. Consequently, a weakly sparse graph class has bounded flip-width if and only if it has bounded expansion.

In Section 6 we study the relationship between twin-width and flip-width. We prove that every class of bounded twin-width has bounded flip-width. We also prove that for classes of ordered graphs, bounded flip-width and bounded twin-width are equivalent.

In Section 7 we prove that classes of bounded flip-width are preserved by first-order transductions.

In Section 8 we introduce a definable variant of flip-width, and prove its equivalence with flip-width. As a consequence, we get a slicewise polynomial (XP) approximation algorithm for the flip-width parameters.

In Section 9 we introduce our third main notion, classes of almost bounded flip-width. We prove that they contain all structurally nowhere dense classes, and study their relationship with monadically stable and monadically dependent classes.

Finally, in Section 10, we discuss possible directions of further research, conjectures and questions.

Results marked with $(*)$ are proved in the appendix.

2 Preliminaries

We introduce basic notation in Section 2.1. In Section 2.2 we recount the fundamental notions of Sparsity theory. In Section 2.3 we recall the notion of Vapnik-Chervonenkis dimension of a set system, of a graph, and of a binary relation. In Section 2.4 we recall basic notions from logic (structures, formulas), and a result characterizing nowhere dense classes in terms of the VC-dimension of certain set systems.

2.1 Notation

$\mathbb{N} = \{0, \dots\}$ denotes the set of natural numbers. For two sets A and B their symmetric difference is denoted $A \triangle B := (A - B) \cup (B - A)$. Graphs are finite, undirected and without self-loops. The set of vertices of a graph G is denoted $V(G)$, and the set of edges of G is denoted $E(G)$. An edge joining u and v is denoted uv. In particular, $uv = vu$ and $u \neq v$ for all $uv \in E(G)$. We write $|G|$ for the number of vertices of G. For a vertex v of a graph G define the (open) *neighborhood* of v in G as $N_G(v) := \{u \mid uv \in E(G)\}$, and write $N(v)$ if G is understood from the context. We write $B^r_G(v)$ to denote the set of vertices at distance at most r from v in G.

A graph H is a *subgraph* of G if H is obtained by removing vertices and/or edges from G, and is an *induced subgraph* of G if H is obtained by removing vertices from G, alongside with the edges incident to them. The subgraph of G induced by a set of vertices $X \subseteq V(G)$ is the graph $G[X]$ with vertices X and edges uv such

that $u, v \in X$ and $uv \in E(G)$. For $X, Y \subseteq V(G)$, the bipartite graph *semi-induced* by X and Y in a graph G has parts X and Y and edges xy such that $x \in X, y \in Y$ and $xy \in E(G)$. Note that X and Y need not be disjoint in G; in this case, $G[X, Y]$ contains two copies of every vertex in $X \cap Y$. Two sets X, Y are *complete* in G if $G[X, Y]$ is the complete bipartite graph, and *anti-complete* in G if $G[X, Y]$ has no edges, and *homogeneous* if they are either complete or anti-complete. K_t denotes the complete graph on t vertices, and $K_{s,t}$ denotes the complete bipartite graph with parts of sizes s and t.

A *k-colored graph* is a graph G equipped with a function, assigning a color from $\{1, \ldots, k\}$ to each vertex of G. A class $\widehat{\mathcal{C}}$ of k-colored graphs is a *k-coloring* of a graph class $\mathcal{C}$, if for every colored graph in $\widehat{\mathcal{C}}$, the underlying uncolored graph belongs to $\mathcal{C}$.

A *graph class* $\mathcal{C}$ is a set of graphs. A class $\mathcal{C}$ is *hereditary* if it is closed under taking induced subgraphs. The *hereditary closure* of a class $\mathcal{C}$ is the class of all induced subgraphs of graphs from $\mathcal{C}$. $\mathcal{C}$ is *weakly sparse* if there is some $t \in \mathbb{N}$ such that every $G \in \mathcal{C}$ excludes the biclique $K_{t,t}$ as a subgraph. A *graph parameter* is a function f assigning reals to graphs, which is invariant under graph isomorphism. For a graph class $\mathcal{C}$ and graph parameter f, denote $f(\mathcal{C}) := \sup_{G \in \mathcal{C}} f(G)$, with $f(\mathcal{C}) = \infty$ if f is unbounded on $\mathcal{C}$. Say that f is *bounded in terms* of g if there is a function $\alpha\colon \mathbb{R} \to \mathbb{R}$ such that $f(G) \leqslant \alpha(g(G))$ for all graphs G. And f and g are *functionally equivalent* if each of them is bounded in terms of the other.

2.2 Sparsity theory

We briefly recall the fundamental notions of Sparsity theory. See [NdM12] for more background.

A graph G is *d-degenerate* if there is a total order on the vertices of G such that every vertex has at most d neighbors before it in the order. The *degeneracy* of G is the least d such that G is d-degenerate.

An *exact r-subdivision* of a graph G is the graph obtained by replacing every edge of G by a path of length $r+1$. If every edge is replaced by a path of length at most $r+1$, the resulting graph is called a *$\leqslant$r-subdivision* of G. For a graph G, let $\tilde{\nabla}_r(G)$ denote the maximum average degree, $2|E(H)|/|V(H)|$, of all graphs H whose $\leqslant r$-subdivision is a subgraph of G. Note that $\text{degeneracy}(G) \leqslant \tilde{\nabla}_0(G) \leqslant 2 \cdot \text{degeneracy}(G)$.

Definition 2.1. *A graph class $\mathcal{C}$ has* bounded expansion *if for every $r \geqslant 1$ we have $\tilde{\nabla}_r(\mathcal{C}) < \infty$.*

Example 2.2. Classes of bounded expansion include many well-studied sparse classes: the class of planar graphs, every class of bounded maximum degree, classes of bounded tree-width, classes that excludes a fixed (topological) minor.

The weak coloring and admissibility numbers are two families of graph parameters that generalize the degeneracy number to higher radii $r \geqslant 1$, and are defined as follows. The *r-weak coloring number* of a graph G, denoted $\text{wcol}_r(G)$, is the smallest number k with the following property: There is a total order $<$ on the vertices of G such that for every vertex v, there are at most k vertices w with $w < v$ that are *weakly r-reachable* from v: there is a path of length at most r from v to w, in which w is the $<$-smallest vertex. On the other hand, the *r-admissibility* of a graph G, denoted $\text{adm}_r(G)$, is the smallest number k with the following property: There

is a total order $<$ on $V(G)$ such that for every $v \in V(G)$ one cannot find $k+1$ paths of length at most r that start at v, end in some vertex $w < v$, and such that any two of the $k+1$ paths share only v as a common vertex.

Both parameters are of central importance in Sparsity theory. It is known that $\mathrm{wcol}_r(G)$ is bounded in terms of $\mathrm{adm}_r(G)$, and vice versa [Dvo13, Lemma 6]. More specifically, we have the following.

Fact 2.3. *For all graphs G and $r \geqslant 1$ we have*

$$\mathrm{adm}_r(G) \leqslant \mathrm{wcol}_r(G) \leqslant O(\mathrm{adm}_r(G))^r. \tag{1}$$

The fundamental notion of Sparsity theory can be characterized using weak coloring numbers as follows:

Fact 2.4 ([Zhu09]). *A class $\mathcal{C}$ of graphs has bounded expansion if and only if $\mathrm{wcol}_r(\mathcal{C}) < \infty$ for every $r \in \mathbb{N}$.*

By (1), we could replace $\mathrm{wcol}_r(\mathcal{C})$ by $\mathrm{adm}_r(\mathcal{C})$ in this characterization.

We use the following inequality.

Fact 2.5. *For all graphs G and $r \geqslant 1$ we have*

$$\mathrm{adm}_r(G) \leqslant O(\tilde{\nabla}_{r-1}(G))^3. \tag{2}$$

We now move to nowhere dense classes.

Definition 2.6. *A graph class $\mathcal{C}$ is* nowhere dense *if for every $r \geqslant 1$ there is some $n \geqslant 1$ such that for all $G \in \mathcal{C}$, no $\leqslant$r-subdivision of K_n is contained as a subgraph of $G \in \mathcal{C}$.*

It is immediate that every class with bounded expansion is nowhere dense, and there are known examples of nowhere dense classes which do not have bounded expansion.

One of the central results of Sparsity theory is the following characterization of nowhere dense classes (recall that $|G|$ is the number of vertices of G).

Fact 2.7. *A hereditary graph class $\mathcal{C}$ is nowhere dense if and only if for every $r \geqslant 1$ and $\varepsilon > 0$ there is a constant $n_{r,\varepsilon}$ such that* $\mathrm{wcol}_r(G) < |G|^\varepsilon$ *for every $G \in \mathcal{C}$ with $|G| > n_{r,\varepsilon}$.*

The condition in Fact 2.7 can be equivalently phrased as follows: for every $r \geqslant 1$ and $\varepsilon > 0$ there is a constant $n_{r,\varepsilon}$ such that $\mathrm{wcol}_r(G) < |G|^\varepsilon$ for every $G \in \mathcal{C}$ with $|G| > n_{r,\varepsilon}$. This can be written more concisely as $\mathrm{wcol}_r(G) \leqslant |G|^{o(1)}$, where $o(1)$ signifies a function $g_r \colon \mathbb{N} \to \mathbb{R}$ (depending on r) with $\lim_{n\to\infty} g_r(n) = 0$.

This fundamental result opens the door for multiple algorithmic applications of nowhere denseness, thanks to the existence of efficient algorithms approximating weak coloring numbers. In particular, the model-checking result [GKS17b] relies on Fact 2.7.

2.3 Vapnik-Chervonenkis dimension

An important parameter measuring the complexity of graphs, and more generally, of set systems, is the *Vapnik-Chervonenkis dimension*, or *VC-dimension*.

A *set system* is a pair $(X, \mathcal{F})$ with $\mathcal{F} \subseteq 2^X$. Its *VC-dimension* is the maximal size of a subset $Y \subseteq X$ such that $\{Y \cap F \mid F \in \mathcal{F}\} = 2^Y$.We recall the fundamental Sauer-Shelah-Perles lemma [Sau72, She72].

Lemma 2.8 (Sauer-Shelah-Perles lemma)**.** *Let $(X, \mathcal{F})$ be a set system of VC-dimension d. Then $|\mathcal{F}| = O(|X|^d)$.*

The VC-dimension of a graph G, denoted $\mathrm{VCdim}(G)$, is defined as the VC-dimension of the set system $(V(G), \{N(v) \mid v \in V(G)\})$. More explicitly, $\mathrm{VCdim}(G)$ is the maximal size of a subset $X \subseteq V(G)$ such that $\{N(v) \cap X \mid v \in V(G)\} = 2^X$.

For a binary relation $R \subseteq V \times W$, an element $v \in V$ and set $S \subseteq W$ denote $R(v; T) := \{s \in S \mid (v,s) \in R\}$ and for $w \in W$ and $S \subseteq V$ let $R(S; w) := \{s \in S \mid (s,w) \in R\}$. The *VC-dimension* of R is the maximum of the VC-dimensions of the two set systems

$$(V, \{R(v; W) \mid v \in V\}) \quad \text{and} \quad (W, \{R(V; w) \mid w \in W\}).$$

2.4 Logic

In this paper, we only consider *binary* signatures, that is signatures Σ consisting of unary relation symbols, binary relation symbols, and unary function symbols. Fix a binary signature Σ. A Σ-structure B consists of a set of vertices $V(B)$, and is equipped with interpretations for each of the symbols in Σ, as unary relations, binary relations, and a unary functions on $V(B)$, respectively. The number of vertices of B is denoted $|B|$. A graph G is seen as a structure over the signature Σ consisting of one binary relation E, which is interpreted as adjacency in G.

Terms are defined inductively, as either a variable, or a function symbol applied to a term. A *quantifier-free* formula $\varphi(x,y)$ over the signature Σ is a boolean combination of *atomic formulas* of the form $U(t(x))$, or $t(x) = t'(y)$, or $R(t(x), t'(y))$, where $U \in \Sigma$ is a unary relation, $R \in \Sigma$ is a binary relation, and $t(x)$ and $t'(y)$ are terms.

A *first-order formula* is built inductively: every atomic formula is a first-order formula, and if φ, ψ are first-order formulas and x is a variable, then so are $\varphi \vee \psi$, $\varphi \wedge \psi$, $\neg\varphi$, $\exists x.\varphi$ and $\forall x.\varphi$. The *quantifier rank* of a formula φ is the maximal nesting of quantifiers in it.

Say that a Σ-formula $\varphi(x,y)$ is *symmetric* if $G \models \varphi(u,v) \leftrightarrow \varphi(v,u)$ holds for every Σ-structure G and elements u, v of G.

The *Gaifman graph* of a Σ-structure B is the graph with vertices $V(B)$ and edges uv such that $(u,v) \in R$ or $(v,u) \in R$ for some binary relation $R \in \Sigma$, or $f(u) = v$ or $f(v) = u$ for some unary function $f \in \Sigma$.

Dependence Let $\varphi(\bar{x}; \bar{y})$ be a first-order formula, whose set of free variables is partitioned into two disjoint sets $\bar{x}$ and $\bar{y}$. For a graph H define the binary relation $R^{\varphi}_{H} \subseteq V(H)^{\bar{x}} \times V(H)^{\bar{y}}$ as

$$R^{\varphi}_{H} = \{(\bar{u}, \bar{v}) \in V(H)^{\bar{x}} \times V(H)^{\bar{y}} \mid H \models \varphi(\bar{u}; \bar{v}).\}.$$

Say that a class $\mathcal{C}$ of Σ-structures is *dependent*, or *NIP* [SS71, Adl08] if for every first-order formula $\varphi(\bar{x}; \bar{y})$ there is $k \geqslant 1$ such that for every $H \in \mathcal{C}$ the binary relation R^{φ}_{H} has VC-dimension at most k. This is equivalent to saying that for every first-order formula $\varphi(\bar{x}; \bar{y})$, there is some bipartite graph G_{φ}, such that for all $H \in \mathcal{C}$, the bipartite graph corresponding to the binary relation R^{φ}_{H} does not contain G_{φ} as an induced subgraph.

The following fact follows from a result of Podewski and Ziegler [PZ78] (see also [AA14] and [PST18]).

Fact 2.9. *Every nowhere dense graph class $\mathcal{C}$ is dependent. Conversely, every monotone, dependent graph class $\mathcal{C}$ is nowhere dense.*

Hereditary, dependent graph classes are much more general than nowhere dense classes. The study of those classes is the main motivation of this paper. A closely related notion, of a *monadically dependent* class, is discussed in Section 9. Monadically dependent classes are dependent, but in general, the converse implication does not hold. However, as shown by Braunfeld and Laskowski [BL22], for hereditary classes, the two notions coincide.

3 Cop-width

We start with defining and analyzing the cop-width game, and the related cop-width parameters. To define the game, we invoke the original description of the Cops and Robber game by Seymour and Thomas [ST93]: "The robber stands on a vertex of the graph, and can at any time run at great speed to any other vertex along a path of the graph. He is not permitted to run through a cop, however. There are k cops, each of whom at any time either stands on a vertex or is in a helicopter (that is, is temporarily removed from the game). The objective of the player controlling the movement of the cops is to land a cop via helicopters on the vertex occupied by the robber, and the robber's objective is to elude capture. (The point of the helicopters is that cops are not constrained to move along paths of the graph – they move from vertex to vertex arbitrarily.) The robber can see the helicopter approaching its landing spot and may run to a new vertex before the helicopter actually lands."

Seymour and Thomas proved that the least number of cops needed to catch a robber on a graph G is equal to one plus the treewidth of G. To this end, they proved a *min-max theorem*: either the cops have a winning strategy of a particularly simple, *monotone* form, which can be described by a tree decomposition of the graph, or otherwise, the robber has a winning strategy of a particularly simple form, called a *haven*.

In our variant of the game the robber runs at speed r, for some fixed $r \in \mathbb{N} \cup \{\infty\}$. That is, in each round, after the cops have taken off in their helicopters to their new positions, which are known to the robber, and before the helicopters have landed, the robber may traverse a path of length at most r that does not run through a cop that remains on the ground (he may also stay put). We call this game the *cop-width game* with radius r and width k, if there are k cops, and the robber can run at speed r.

Definition 3.1. *The* radius-r cop-width *of G, denoted* $\mathrm{copwidth}_r(G)$, *is the least number k such that the cops win the cop-width game with radius r and width k.*

Note that there is a graph parameter called the *cop-number* of a graph [AF84], which is equal to the number of cops needed to catch the robber in a game where the cops and the robber move at speed one in each turn. It could seem that $\mathrm{copwidth}_1(G)$ is upper bounded by the cop-number of G. However, there is a crucial difference with our notion: the cops do not announce their moves in advance. So for instance, every graph with a universal vertex (a vertex adjacent to

all other vertices) has cop-number equal to one (such graphs are called *cop-win graphs*). On the other hand, $\mathrm{copwidth}_1(G)$ can be arbitrarily large on such graphs. Indeed, it is easy to see that for each $r \in \mathbb{N} \cup \{\infty\}$, the parameter $\mathrm{copwidth}_r(G)$ is monotone with respect to the subgraph relation: if H is a subgraph of G then $\mathrm{copwidth}_r(H) \leqslant \mathrm{copwidth}_r(G)$.

Example 3.2. Let G be a graph with maximum degree d. Then $\mathrm{copwidth}_r(G) < d^{r+1}$, for all $r \in \mathbb{N}$. To see this, consider the following strategy for the cops: if the robber is initially placed on a vertex v, direct the cops to all (less than d^{r+1}) vertices that are at distance at most r from v. Before the cops land at those locations, the robber traverses a path of length at most v to one of those vertices, and is caught by the landing cop.

Example 3.3. Let T be a rooted tree. Then G has tree-width at most one, so by the Seymour-Thomas result, $\mathrm{copwidth}_\infty(G) \leqslant 2$. The strategy with two cops is as follows. In the first round, land one cop on the root of the tree, leaving the other cop in his helicopter. In each subsequent round, direct the cop that is further from the robber, to the child of the other cop's node which is an ancestor of (or equal to) the robber's node. With this strategy, in round i the robber will be at distance at least i from the root, so he will be caught after at most as many rounds as the height of T.

We start with comparing the most fundamental parameter of Sparsity theory, degeneracy, with the cop-width parameter for radius 1.

Theorem 3.4. *For every graph G,*

$$\mathrm{copwidth}_1(G) = \mathrm{degeneracy}(G) + 1.$$

The simple proof of this fact relies on the fundamental duality result concerning degeneracy: every graph G is either d-degenerate, or it has a subgraph H in which every vertex has degree larger than d. This duality theorem is nothing else than a min-max theorem for the radius-1 cop game, quite analogous to the min-max theorem in the eponymous paper of Seymour and Thomas. Indeed, a d-degeneracy order can now be viewed as a compact representation of a winning strategy for the cops involving $d+1$ cops, in the cop-width game with radius 1: when robber is on a vertex v, place $d+1$ cops on v and the neighbors of v before v. Then the robber needs to move rightwards in the order, and eventually loses. This proves $\mathrm{copwidth}_1(G) \leqslant \mathrm{degeneracy}(G) + 1$. Dually, a subgraph H of G whose all vertices have degree larger than d, can be seen as a haven for the robber: he can forever evade d-cops by always moving to an unoccupied vertex of H (or remaining in place). This proves $\mathrm{copwidth}_1(G) \geqslant \mathrm{degeneracy}(G) + 1$, thus proving Theorem 3.4.

Next we observe that for higher radii r, the parameter $\mathrm{copwidth}_r(G)$ is closely related to the generalized coloring numbers: the weak coloring number $\mathrm{wcol}_r(G)$ and the admissibility numbers $\mathrm{adm}_r(G)$ (see Section 2.2). Recall (see Fact 2.3) that the two parameters are functionally equivalent.

We prove the following:

Theorem 3.5 ($*$). *For $r \in \mathbb{N}$,*

$$\mathrm{adm}_r(G) + 1 \leqslant \mathrm{copwidth}_r(G) \leqslant \mathrm{wcol}_{2r}(G) + 1.$$

In particular, by Fact 2.4, this gives the first, arguably very natural, characterization of classes with bounded expansion, in terms of a game:

Corollary 3.6. *A graph class $\mathcal{C}$ has bounded expansion if and only if for every $r \in \mathbb{N}$ we have that* $\mathrm{copwidth}_r(\mathcal{C}) < \infty$*.*

Thus, all the classes $\mathcal{C}$ from Example 2.2 satisfy $\mathrm{copwidth}_r(\mathcal{C}) < \infty$, for all $r \in \mathbb{N}$. Similarly, by Fact 2.7 we get a new characterization of nowhere dense classes.

Corollary 3.7. *A hereditary graph class $\mathcal{C}$ is nowhere dense if and only if for every $r \in \mathbb{N}$ and $\varepsilon > 0$ we have that* $\mathrm{copwidth}_r(G) = O_{r,\varepsilon}(|G|^{\varepsilon})$ *for $G \in \mathcal{C}$.*

The proof of Theorem 3.5, which we now give, is quite simple thanks to the theory of sparsity, but also sheds a new light on the fundamental notions of that theory. The total order $<$ that appears in the definition of the weak coloring number $\mathrm{wcol}_{2r}(G)$ can be viewed as a compact representation of a (very particular) winning strategy for the cops in the cop-width game with radius r. Indeed, the following yields a winning strategy for the cops: if the robber is at a vertex v, then the cops are placed on v and the vertices $w < v$ that are $2r$-weakly reachable from v. To see that this strategy is winning for the cops, consider the paths $\pi_1, \pi_2, \ldots$ in G, where π_i is the path of length at most r along which robber traversed from his ith position v_i to his $(i+1)$-st position v_{i+1}. If m_i denotes the $<$-minimal vertex of the path m_i, then we have that $m_{i+1} > m_i$. Otherwise, m_{i+1} is $2r$-weakly reachable from v_i, so at the time the robber traversed the path π_{i+1}, the vertex m_{i+1} that was occupied by a cop, which is impossible. Therefore, we have $m_1 < m_2 < \ldots$, so the cops win after at most $|G|$ rounds. This gives the upper bound in Theorem 3.5.

For the lower bound, we use the well-known (and straightforward) min-max characterization of admissibility: a graph G has r-admissibility number at least d if there is a set of vertices $X \subseteq V(G)$ such that for every vertex $v \in X$ there is a set of d paths of length at most r that start at v and end in some vertex of $X - v$, such that any two paths only share v as a common vertex.

A set $X \subseteq V(G)$ that witnesses that $\mathrm{adm}_r(G) \geqslant d$ can used as a haven for the robber, to elude d cops forever, similarly as in the case of degeneracy. If the robber is occupying a vertex v in X and the cops are moving to a set S of at most d new positions, then either $v \notin X$, or there is some path from v to a vertex in $X - S$ of length r, and the robber moves along this path. This proves $\mathrm{copwidth}_r(G) \geqslant d+1$.

Therefore, the equivalence of the weak coloring numbers and the admissibility numbers can again be seen as a min-max theorem for the cop-width game with radius r.

Note that Theorem 3.5 does not give an exact min-max theorem, as there is a gap between the upper and lower bounds. We can get a family of parameters based on another variant of the cops and robber game, which does admit an exact min-max theorem. In this game, in each round first the cops move to some k vertices of the graph, and then the robber moves via a path of length at most r, that does not run through a cop, and loses if no such path exists. See Appendix A.1 for more details. This family of parameters again characterizes classes with bounded expansion (see Corollary A.3), and does admit an exact min-max theorem (Lemma A.1) that exhibits a duality between total orders describing winning strategies for the cops, and havens for the robber. Those parameters essen-

tially[5] appear in the work [RT08, LPPT20]. However, the authors fail to make the connection with generalized coloring numbers, and with classes of bounded expansion. Curiously, the limit version of those parameters, for radius $r = \infty$, does not correspond to treewidth, but to a notion called ∞-*admissibility* [Dvo12]. Classes with bounded ∞-admissibility are characterized as clique-sums of graphs of almost bounded degree [Dvo12, Cor. 5 and Thm. 6]. Despite the appealing properties of the parameters based on this variant of the cops and robber game, they seem to be less suited for our purposes, of generalizing to dense graphs.

To summarize, our new cop-width parameters exactly characterize treewidth (for $r = \infty$), degeneracy (for $r = 1$), classes of bounded expansion, and nowhere dense classes. This captures an appreciable fragment of the theory of sparse graphs, while offering a new perspective on the fundamental graph parameters used for measuring sparsity, and the dualities between them. We have

$$\begin{aligned}\mathrm{degeneracy}(G) + 1 = \mathrm{copwidth}_1(G) \leqslant \mathrm{copwidth}_2(G) \leqslant \ldots\\ \leqslant \mathrm{copwidth}_\infty(G) = \mathrm{treewidth}(G) + 1,\end{aligned}$$

so for a class $\mathcal{C}$ of graphs, if any of those parameters is bounded by a constant, then $\mathcal{C}$ has bounded degeneracy by Theorem 3.4. So those parameters are only well suited to the study of sparse graphs: every class $\mathcal{C}$ for which either of those parameters is bounded, is sparse, in the sense of having a bound on the edge density $|E(G)|/|V(G)|$ for all graphs G in the class. Moreover, $\mathrm{copwidth}_r$ is monotone with respect to the subgraph relation: if H is a subgraph of G, then $\mathrm{copwidth}_r(H) \leqslant \mathrm{copwidth}_r(G)$.

4 Flip-width

To lift the cop-width game to the setting of dense graphs, we enhance the power of the cops. Now, instead of placing cops on at most k vertices of the graph, which can be alternatively seen as removing at most k vertices, or isolating them, the player controlling the cops can perform *flips* on subsets of the graph G. For a fixed graph G, applying a *flip* between a pair of sets of vertices $A, B \subseteq V(G)$ results in the graph obtained from G by inverting the adjacency between any pair of vertices a, b with $a \in A$ and $b \in B$. For example, applying a flip between $V(G)$ and $V(G)$ in G results in the complement of G. And if v is a vertex of G, then applying a flip between $\{v\}$ and the neighborhood $N(v)$ of v has the same effect as *isolating* v, that is, removing all the edges adjacent to v in G. If G is a graph and $\mathcal{P}$ is a partition of its vertex set, then call a graph G' a $\mathcal{P}$-*flip* of G if G' can be obtained from G by performing a sequence of flips between pairs of parts $A, B \in \mathcal{P}$ (possibly with $A = B$). Since flips are involutive and commute with each other, such a sequence of flips can be specified by a set of at most $\binom{|\mathcal{P}|+1}{2}$ unordered pairs of elements of $\mathcal{P}$. Finally, call G' a k-*flip* of G, if G' is a $\mathcal{P}$-flip of G, for some partition $\mathcal{P}$ of $V(G)$ with $|\mathcal{P}| \leqslant k$.

We remark that there are many other, functionally equivalent, ways to measure the complexity of a k-flip (also called a *perturbation*) G' of G. For example, say that

[5] The paper [RT08] considers a variant of the game in which the robber is *lazy*, that is, does not move unless a cop is placed at his location, whereas the [LPPT20] considers a variant where the cops occupy edges instead of vertices, and the robber never remains put.

G' is a *k-sequential-flip* of G if G' is obtained from G by applying a sequence of flips between k pairs of arbitrary subsets of $V(G)$. If G' is a k-sequential-flip of G then G' is a 2^{2k}-flip of G, and conversely, if G' is a k-flip of G, then G' is a $\binom{k+1}{2}$-sequential-flip of G. We could also require that flips are only applied to pairs of the form (A, A); this would lead to a functionally equivalent parameter, as flipping a pair A, B can be obtained by flipping three pairs: $(A \cup B, A \cup B)$, (A, A) and (B, B). Other, functionally equivalent measures of the complexity of a flip G' of a graph G can be defined by considering the graph $G' \triangle G$ with vertices $V(G)$ and edges $E(G') \triangle E(G)$. Note that G' is a k-flip of G if and only if $G' \triangle G$ is a k-flip of the edgeless graph on $V(G)$. This is equivalent to $G \triangle G'$ having *neighborhood diversity* [Lam12] k. The rank of the adjacency matrix of $G \triangle G'$ over a fixed finite field, see [NiO20, DK06] leads to a further, functionally equivalent complexity measure of a flip.

We now define our central notions. The *flip-width* game with radius $r \in \mathbb{N} \cup \{\infty\}$ and width $k \in \mathbb{N}$, $k \geqslant 1$, is played on a graph G. In each round i of the game, a k-flip G_i of G is declared by the cops, and the new position $v_i \in V(G)$ is selected by the robber, as follows. Initially, $G_0 = G$ and v_0 is a vertex of G chosen by the robber. In round $i > 0$, the cops announce a new k-flip G_i of G, that will be put into effect momentarily. The robber, knowing G_i, moves to a new vertex v_i by following a path of length at most r from v_{i-1} to v_i in the *previous* graph G_{i-1}. The game terminates when the robber is trapped, that is, when v_i is isolated in G_i.

Definition 4.1. *Fix $r \in \mathbb{N} \cup \{\infty\}$. The* radius-$r$ flip-width *of a graph G, denoted $\mathrm{fw}_r(G)$, is the smallest number $k \in \mathbb{N}$ such that the cops have a winning strategy in the flip-width game of radius r and width k on G.*

Definition 4.2. *A class $\mathcal{C}$ of graphs has* bounded flip-width *if $\mathrm{fw}_r(\mathcal{C}) < \infty$ for every $r \in \mathbb{N}$. More explicitly: for every radius $r \in \mathbb{N}$ there is some $c_r \in \mathbb{N}$ such that $\mathrm{fw}_r(G) < c_r$ for all $G \in \mathcal{C}$.*

We remark that the cop-width parameters considered in Section 3 are functionally equivalent correspond (more precisely, each parameter can be bounded from above by a linear function of the other) to parameters defined by a variant of the flip-width game, call it the *isolation game*, which is played as the flip-width game, but each graph G_i announced by the cops is obtained from G by isolating at most k vertices in G (see Lemma A.4). The difference between the isolation game and the cop-width game is that in the cop-width game, the robber can move through a vertex from which a cop has just departed by a helicopter, while in the isolation game, he cannot.

We now argue that the flip-width parameters are bounded in terms of the corresponding cop-width parameters.

Lemma 4.3. *For every $r \in \mathbb{N} \cup \{\infty\}$ and graph G, we have*

$$\mathrm{fw}_r(G) \leqslant \mathrm{copwidth}_r(G) + 2^{\mathrm{copwidth}_r(G)}. \tag{3}$$

Proof. The main observation is that isolating a set S of at most k vertices in G can be achieved by performing a $(k + 2^k)$-flip: consider the partition $\mathcal{P}_S$ that partitions S into singletons and $V(G) - S$ according to the neighborhood in S, and flip $\{s\}$ with every class of the partition that is complete to $\{s\}$. Note that $|\mathcal{P}_S| \leqslant k + 2^k$. Now, if the cops have a winning strategy in the cop-width game of radius r and

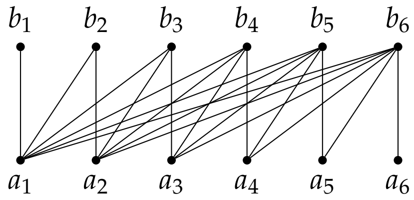

Figure 1: A half-graph of order 6.

width k on a graph G, we can use this strategy in the flip-width game of radius r and width $k+2^k$, as follows: whenever the cops announce a new set S of positions of the cops in the cop-width game, in the flip-width game, the cops announce the graph G' with the vertices in S isolated, which is a $(k+2^k)$-flip of G. It is easy to verify that if the cops win in the cop-width game, then, playing according to the above strategy, they also win in the flip-width game. Inequality (3) follows. $\square$

The following gives an improved bound. Note that if $\mathrm{copwidth}_1(G) \leqslant t$ then G excludes $K_{t,t}$ as a subgraph, by Theorem 3.4 and the fact that $K_{t,t}$ is not $(t-1)$-degenerate. From this, one can bound the size of the partition $\mathcal{P}_S$ considered above by k^t (see Lemma B.3), and obtain:

Theorem 4.4. *Fix $r \in \mathbb{N} \cup \{\infty\}$. Let G be a graph and let t be the smallest number such that G excludes $K_{t,t}$ as a subgraph; in particular $t \leqslant \mathrm{copwidth}_1(G) \leqslant \mathrm{copwidth}_r(G)$. Then*

$$\mathrm{fw}_r(G) \leqslant \mathrm{copwidth}_r(G)^t.$$

In particular, every class $\mathcal{C}$ with bounded expansion has bounded flip-width.

Whereas the cop-width parameters are monotone with respect to the subgraph relation, the flip-width parameters are monotone with respect to the induced subgraph relation.

Lemma 4.5. *Fix $r \in \mathbb{N} \cup \{\infty\}$. If H is an induced subgraph of G, then $\mathrm{fw}_r(H) \leqslant \mathrm{fw}_r(G)$. In particular, a class $\mathcal{C}$ has bounded flip-width if and only if its hereditary closure has bounded flip-width.*

4.1 Examples

We start by giving some example classes of bounded flip-width. By Theorem 4.4, every class of bounded expansion has bounded flip-width, see Example 2.2 for some specific classes. Unlike cop-width, flip-width is not limited to sparse graphs, and is geared towards the study dense graphs.

Example 4.6. If $\bar{G}$ is the complement of G and $r \in \mathbb{N} \cup \{\infty\}$, then $\mathrm{fw}_r(\bar{G}) = \mathrm{fw}_r(G)$, since any k-flip of G is also a k-flip of G'. Therefore, if $\mathcal{C}$ is a class with bounded expansion, then the class $\bar{\mathcal{C}} := \{\bar{G} \mid G \in \mathcal{C}\}$, has bounded flip-width, and if $\mathcal{C}$ has bounded treewidth, then $\mathrm{fw}_\infty(\bar{\mathcal{C}}) = \mathrm{fw}_\infty(\mathcal{C}) < \infty$. In particular, as every edgeless graph G has $\mathrm{fw}_\infty(G) = 1$, it follows that every clique G also has $\mathrm{fw}_\infty(G) = 1$.

Example 4.7. Consider the *half-graph* H_n of order n, as depicted in Figure 1. We show that $\mathrm{fw}_\infty(H_n) \leqslant 4$. Observe that applying a flip between $\{a_1, \ldots, a_{i-1}\}$

and $\{b_i, \ldots, b_n\}$ breaks the half-graph into two connected components. Additionally flipping $\{a_i\}$ and $\{b_i, \ldots, b_n\}$ makes a_i and b_i isolated. The strategy of the cops is to perform the above two flips in their ith move, thus pushing the robber rightwards in each round (to perform those flips, partition $V(H_n)$ into four parts: $\{1, \ldots, a_{i-1}\}, \{a_i\}, \{b_i, \ldots, b_n\}$, and the rest).

Example 4.8. The comparability graph G of a rooted tree T is the graph with vertices $V(T)$, which are adjacent if and only if one is an ancestor of the other in T. Generalizing half-graphs, those graphs also have $\mathrm{fw}_\infty(G) \leqslant 4$. The strategy of the cops is, in round i, to consider the node u_i at depth i that is the ancestor of the current position of the robber, and to isolate u_i, and remove all edges between the descendants of u_i and the ancestors of u_i. This can be done by partitioning $V(G)$ into four parts: $\{u_i\}$, the ancestors of u_i, the descendants of u_i, and the rest.

Example 4.9. Fix $r \in \mathbb{N} \cup \{\infty\}$. If $G_1, \ldots, G_m$ are graphs and G is their disjoint union, then the following inequality holds:

$$\mathrm{fw}_r(G) \leqslant \max_{1 \leqslant i \leqslant m} (\mathrm{fw}_r(G_i)) + 1.$$

The $+1$ comes from the fact that a partition $\mathcal{P}$ of $V(G_i)$ into k parts induces a partition of $V(G)$ into $k+1$ parts, namely the k parts of $\mathcal{P}$, and $V(G) - V(G_i)$. Thus, if the robber is hiding in the graph G_i, the cops may translate a winning strategy on G_i of width $k := \mathrm{fw}_r(G_i)$, into a winning strategy on G of width $k+1$.

It follows that if $\mathcal{C}$ has bounded flip-width, then the class of disjoint unions of graphs from $\mathcal{C}$ has bounded flip-width.

A *modular partition* of a graph G is a partition $\mathcal{P}$ of $V(G)$ such that any two distinct parts are homogeneous in G. The quotient graph $G/\mathcal{P}$ has as vertices the parts of $\mathcal{P}$, and as edges pairs of distinct parts that are complete in G. An extension of the idea in Example 4.9 yields the following.

Lemma 4.10 ($*$). *Let G be a graph and $\mathcal{P}$ be its modular partition. Then*

$$\mathrm{fw}_r(G) \leqslant \max\left(\mathrm{fw}_r(G/\mathcal{P}), \max_{A \in \mathcal{P}} \mathrm{fw}_r(G[A]) + 2\right).$$

The strategy for the cops on G first follows the strategy on $G/\mathcal{P}$, where each k-flip of $G/\mathcal{P}$ is lifted naturally to a k-flip of G. Once a part $A \in \mathcal{P}$ is isolated in the game on $G/\mathcal{P}$, the strategy on $G[A]$ is used. The $+2$ in the statement is due to the fact that every partition $\mathcal{Q}$ of A into k parts induces a partition of G into $k+2$ parts: the k parts of $\mathcal{Q}$, the (common) set of neighbors of vertices in A outside of A, and the rest. See Appendix C for details. In particular, if G is the lexicographic product of two graphs H and K (obtained by blowing up each node of H to a copy of K) then $\mathrm{fw}_r(G) \leqslant \max(\mathrm{fw}_r(H), \mathrm{fw}_r(K)+2)$.

The *substitution closure* of a class of graphs $\mathcal{C}$ is the smallest class $\mathcal{C}^*$ containing $\mathcal{C}$ such that if G is a graph with a modular partition $\mathcal{P}$ into modules A satisfying $G[A] \in \mathcal{C}^*$, and $G/\mathcal{P} \in \mathcal{C}$, then $G \in \mathcal{C}^*$. Intuitively, a graph in $\mathcal{C}^*$ can be obtained from a single vertex by repeatedly blowing up vertices to graphs from $\mathcal{C}$.

Using similar ideas as in Lemma 4.10, we prove Lemma 4.11.

Lemma 4.11 ($*$). *For every $r \in \mathbb{N} \cup \{\infty\}$ and graph class $\mathcal{C}$, we have*

$$\mathrm{fw}_r(\mathcal{C}^*) \leqslant \mathrm{fw}_r(\mathcal{C}) + 2.$$

In particular, if $\mathcal{C}$ has bounded flip-width, then $\mathcal{C}^*$ has bounded flip-width.

Example 4.12. The substitution closure $\mathcal{C}^*$ of the class $\mathcal{C}$ of subcubic graphs has bounded flip-width. This follows from Example 3.2, Theorem 4.4, and Lemma 4.11.

More examples are given in the following sections: they include classes of bounded clique-width (see Section 4.2), classes of bounded twin-width (see Section 6), and interpretations of classes of bounded expansion (see Section 7). We remark that the class from Example 4.12 is not among those classes (it has unbounded twin-width, and is not edge-stable; see Section 9).

4.2 Flip-width with infinite radius

We now focus on the parameter fw_∞. Generalizing Example 4.8, we observe that any class of bounded clique-width, equivalently, of bounded rank-width, has bounded fw_∞. We briefly sketch the argument now. If G has clique-width at most k, then there is a tree T with leaves $V(G)$ such that for every edge e of T, if $V(G) = X \uplus Y$ is the bi-partition induced by e (into the leaves on either side of e), then there is a partition $X = X_1 \uplus \cdots \uplus X_p$ and a partition $Y = Y_1 \uplus \cdots \uplus Y_q$ with $p, q = O(k)$, and such that X_i and Y_j are complete in G. This implies that there is a $O(k^2)$-flip of G which removes all the edges with one endpoint in X and one endpoint in Y. This can be used by the cops in their winning strategy, to restrain the robber to the leaves of smaller and smaller subtrees of T, similarly as in Example 4.8. See Appendix D for details.

In fact, similarly as $\mathrm{copwidth}_\infty$ is equivalent to tree-width, we show that fw_∞ is functionally equivalent to the clique-width, or rank-width parameters. To the best of our knowledge, this gives the first characterization of graph classes of bounded clique-width, in terms of a game[6].

Theorem 4.13 ($*$). *A class of graphs $\mathcal{C}$ has bounded clique-width if and only if $\mathrm{fw}_\infty(\mathcal{C}) < \infty$.*

The easier, forward implication is sketched above (see Appendix D for more details). The backwards implication is more involved, as it relies on nontrivial results describing obstructions that can be found in classes of unbounded clique-width.

We give a quick proof of the backwards implication, using two results that are shown in Section 7 and Section 4.4, and a result of Courcelle and Oum. See Section 7 for a definition of a CMSO transductions. Courcelle and Oum [CiO07, Corollary 7.5] proved the following result[7].

Theorem 4.14. *Every graph class of unbounded clique-width* CMSO-*transduces the class of all graphs.*

On the other hand, in Section 7, we prove the following.

Corollary (7.7)**.** *If $\mathcal{C}$ is a class with $\mathrm{fw}_\infty(\mathcal{C}) < \infty$ that* CMSO-*transduces transduces a class $\mathcal{D}$, then $\mathrm{fw}_\infty(\mathcal{D}) < \infty$.*

In Section 4.4, Corollary 4.20 we prove that there are graphs G with arbitrarily large $\mathrm{fw}_1(G)$. In particular, if $\mathcal{D}$ is the class of all graphs, then $\mathrm{fw}_\infty(\mathcal{D}) = \infty$. Those

[6] The radius-∞ flip-width game arose in a private discussion in 2018 with Michał Pilipczuk.

[7] In fact, they proved the same statements, but with 'all square grids' instead of 'all graphs'. However, it is well known and easy to see that the class of all square grids CMSO-transduces the class of all graphs, and transductions can be composed

results give the backward implication in Theorem 4.13, since by Corollary 7.7, $\mathcal{D}$ does not transduce in any class $\mathcal{C}$ with $\mathrm{fw}_\infty(\mathcal{C}) < \infty$. By Theorem 4.14, every class with $\mathrm{fw}_\infty(\mathcal{C}) < \infty$ has bounded clique-width.

4.3 Hideouts

We move to the study of finite radii, which are our main focus. First, we describe a notion of a hideout, which is a useful way of describing a winning strategy of the robber, thus proving lower bounds on $\mathrm{fw}_r(G)$. In Section 4.4 we use this notion to prove some combinatorial properties of graphs with bounded flip-width.

Definition 4.15. *Fix $r, k, d \geqslant 1$. A (r, k, d)-*hideout *in a graph G is a set of vertices $U \subseteq V(G)$ with $|U| > d$, satisfying the following property. For every k-flip G' of G,*

$$|\left\{v \in U : |B^r_{G'}(v) \cap U| \leqslant d\right\}| \leqslant d,$$

that is, there are at most d vertices $v \in U$ for which there are at most d vertices $u \in U$ that are connected with v by a path of length at most r in G'.

Lemma 4.16. *Fix $r, k \geqslant 1$. Suppose that G has a (r, k, d)-hideout U for some $d \geqslant 1$. Then $\mathrm{fw}_r(G) > k$.*

Proof. Let $U \subseteq V(G)$ be a (r, k, d)-hideout. We describe a strategy for the robber in the flip-width game on G with radius r and width k, which allows him to elude the cops indefinitely. The strategy is as follows: when the cops announce a k-flip G' of G, the robber moves to some vertex $v \in U$ such that $|B^r_{G'}(v) \cap U| > d$. In the first move, pick any $v \in U$ with $|B^r_G(v) \cap U| > d$.

We show it is always possible to make a move as described in the strategy. Suppose at some point in the game, the current position v of the robber is such that

$$|B^r_P(v) \cap U| > d. \tag{4}$$

where P is the previous k-flip of G announced by the cops (in the first round, $P = G$), and that the cops now announce a k-flip N of G. Since U is a (r, k, d)-hideout, the set $X \subseteq U$ of vertices $w \in U$ such that $|B^r_N(w) \cap U| \leqslant d$ satisfies $|X| \leqslant d$. By (4), $B^r_P(v)$ contains at least one vertex $v' \in U - X$. The robber moves from v to v' along a path of length at most r in P. As $v' \in U - X$, the invariant is maintained.

Therefore, playing according to this strategy, the robber can elude the cops indefinitely, so $\mathrm{fw}_r(G) > k$. □

Although we use hideouts on several occasions to prove lower bounds on flip-width, we do not know whether the existence of hideouts is a necessary condition for having large flip-width. This is stated as Question 10.3 in Section 10.

4.4 Radius-one flip-width

We now focus on the first parameter, fw_1. We have seen that its sparse analogue, $\mathrm{copwidth}_1$, corresponds precisely to degeneracy (plus one), which is a very well-understood parameter, with many good algorithmic and combinatorial properties. The parameter fw_1 enjoys many useful combinatorial properties, relating it to near-twins, and to the VC-dimension.

Near-twins We prove a first combinatorial property of graphs with small $\mathrm{fw}_1(G)$, namely that such graphs have near-twins. This has several consequences. Say that two vertices u, v of a graph G are *δ-near-twins* if $|N(u) \triangle N(v)| \leqslant \delta$, where $\triangle$ denotes the symmetric difference. We show that every graph G with $\mathrm{fw}_1(G) \leqslant k$ has a pair of $2k$-near-twins. More generally, we prove:

Lemma 4.17. *Let $b, k \in \mathbb{N}$ and let G be a graph with $\mathrm{fw}_1(G) \leqslant k$. Then G contains a set of at least $b+1$ vertices which are mutual $2bk$-near-twins.*

Proof. Assume that G has no set containing $b+1$ mutual $2bk$-near-twins. We prove that $V(G)$ is a $(1, k, bk)$-hideout in G, which implies that $\mathrm{fw}_1(G) > k$ by Lemma 4.16.

Let G' be a k-flip of G and let B be the set of vertices of degree at most bk in G'. We show that $|B| \leqslant bk$, proving that $V(G)$ is a $(1, k, bk)$-hideout in G.

Suppose that $|B| > bk$. Let $\mathcal{P}$ be the partition defining the k-flip G' of G with $|\mathcal{P}| \leqslant k$. As $|B| > bk$, there is a set $B_0 \subseteq B$ with $|B_0| > b$, such that B_0 is contained in one part of $\mathcal{P}$. Any two vertices of B_0 are $2bk$-near-twins in G', as they both have degree at most bk in G'. Since B_0 is contained in a single part of $\mathcal{P}$, it follows that any two vertices of B_0 are $2bk$-near-twins in G, too. But $|B_0| > b$, so this contradicts the assumption. Hence, $|B| \leqslant bk$. □

A bipartite variant of Lemma 4.17, with a very similar proof, is as follows. Recall that if $X, Y \subseteq V(G)$ are two sets of vertices of a graph G then $G[X, Y]$ denotes the bipartite graph semi-induced by X and Y in G.

Lemma 4.18 ($*$)**.** *Let $b, k \in \mathbb{N}$ and let G be a graph with $\mathrm{fw}_1(G) \leqslant k$. Then for every two sets $X, Y \subseteq V(G)$ there is a set $A \subseteq V(G[X, Y])$ consisting of $b+1$ mutual $2bk$-near-twins, which is contained in one of the two parts of $G[X, Y]$.*

Setting $b := 1$ in Lemma 4.17 we get:

Corollary 4.19. *Let G be a graph with $\mathrm{fw}_1(G) \leqslant k$. Then G has a pair of $2k$-near-twins.*

We can now verify that there exist graphs G with arbitrarily large $\mathrm{fw}_1(G)$. It is well known that there exist graphs of arbitrarily large girth and minimum degree.

Corollary 4.20. *A graph G with girth $g \geqslant 4$ and minimum degree larger than k has $\mathrm{fw}_1(G) > k$. Therefore, there exist graphs G with arbitrarily large $\mathrm{fw}_1(G)$.*

Proof. Any two distinct vertices u and v have at most one common neighbor, so $|N(u) \triangle N(v)| > 2k$, so G has no pair of $2k$-near-twins, hence $\mathrm{fw}_1(G) > k$ by Corollary 4.19. □

VC dimension We now show that the VC-dimension of a graph G is bounded in terms of $\mathrm{fw}_1(G)$. Set systems and graphs of bounded VC-dimension have many useful properties, and we use one of them later for studying classes of bounded flip-width. We also consider a related parameter, called *2VC-dimension* [BT15], and denoted $2\mathrm{VCdim}(G)$. This is the maximal size of a set $X \subseteq V(G)$ such that for every two distinct $a, b \in X$ there is a vertex $c \in V(G)$ with $N_G(c) \cap X = \{a, b\}$. Clearly, $\mathrm{VCdim}(G) \leqslant 2\mathrm{VCdim}(G)$. We prove the following.

Theorem 4.21 ($*$)**.** *For every graph G we have*

$$\mathrm{VCdim}(G) \leqslant 8\,\mathrm{fw}_1(G), \tag{5}$$
$$2\mathrm{VCdim}(G) \leqslant 8\,\mathrm{fw}_2(G) + 2 \tag{6}$$

Note that 2VCdim cannot be bounded in terms of fw_1, only in terms of fw_2, as witnessed by 1-subdivided cliques, which are 2-degenerate, and hence (by Theorem 4.4 and Theorem 3.4), have bounded fw_1, and clearly have unbounded fw_2. Furthermore, graphs of girth larger than 4 have VC-dimension at most two, but have arbitrarily large fw_1 by Corollary 4.20. Hence, fw_1 is not bounded in terms of $\mathrm{VCdim}(G)$.

Inequality (5) in Theorem 4.21 follows from Lemma 4.18 (for $b = 1$) and the following.

Lemma 4.22 (∗)**.** *Let G be a graph with $\mathrm{VCdim}(G) \geqslant 2^m$, for some m. Then there are two sets X, Y such that the bipartite graph $G[X, Y]$ contains no pair of $(2^{m-1} - 1)$-near-twins in either of the parts X, Y.*

In the proof, the sets X and Y are two copies of the m-dimensional vector space over the two-element field, with edges connecting vectors with a nonzero dot product. See Appendix E.1.

We prove inequality (6) in Corollary 5.7 later.

From Lemma 4.22 and Lemma 4.18 (for $b = 1$) we conclude the following.

Corollary 4.23. *Let G be a graph with $\mathrm{VCdim}(G) \geqslant d$. Then G contains an induced subgraph H with $O(d)$ vertices and with $\mathrm{fw}_1(H) \geqslant d/8$.*

Corollary 4.24. *If $\mathcal{C}$ is a hereditary class of graphs such that $\mathrm{fw}_1(G) = o(n)$ for every n-vertex graph $G \in \mathcal{C}$, then $\mathrm{VCdim}(\mathcal{C}) < \infty$.*

Proof. If $\mathrm{VCdim}(\mathcal{C}) = \infty$ then for every d there is a graph $G \in \mathcal{C}$ with $\mathrm{VCdim}(G) \geqslant d$, and by Corollary 4.23 there is $H \in \mathcal{C}$ with $O(d)$ vertices and $\mathrm{fw}_1(H) = \Omega(d)$. Since this holds for all $d \in \mathbb{N}$, it cannot be that $\mathrm{fw}_1(G) = o(n)$ for every n-vertex graph $G \in \mathcal{C}$. □

5 Flip-width in weakly sparse classes

We have seen in Theorem 4.4 radius-one flip-width is upper bounded in terms of degeneracy, and radius-r flip-width is upper bounded in terms of generalized coloring numbers. In this section, we provide bounds in the other direction, in weakly sparse classes. It follows that for weakly sparse classes, having bounded flip-width is equivalent to having bounded expansion.

5.1 Radius-one flip-width and degeneracy

As fw_1 is bounded in terms of $\mathrm{copwidth}_1$, which is equivalent to degeneracy by Theorem 3.4, it follows that every class of bounded degeneracy has bounded fw_1. Clearly, every class of degeneracy bounded by t is weakly sparse, as it excludes $K_{t+1,t+1}$ as a subgraph. We show that for weakly sparse classes, bounded degeneracy is equivalent to having bounded fw_1.

Theorem 5.1. *If G is a graph that avoids $K_{t,t}$ as a subgraph, then*

$$\mathrm{degeneracy}(G)/(2t^2) < \mathrm{fw}_1(G) \leqslant (\mathrm{degeneracy}(G) + 1)^t. \tag{7}$$

As a consequence, if $\mathcal{C}$ is a weakly sparse class of graphs then $\mathrm{fw}_1(\mathcal{C}) < \infty$ if and only if $\mathcal{C}$ has bounded degeneracy.

Proof. The second inequality is by Theorem 4.4 and Theorem 3.4. To prove the first inequality, we show that $\text{degeneracy}(G) < 2kt^2$, where $k := \text{fw}_1(G)$.

Recall that G is d-degenerate if and only if every induced subgraph of G contains a vertex of degree at most d. Therefore, to prove $\text{degeneracy}(G) < 2kt^2$ it is enough to show that G contains some vertex of degree less than $2kt^2$ (since the same holds for every induced subgraph H of G, as $\text{fw}_1(H) \leqslant \text{fw}_1(G) \leqslant k$).

Setting $b := t - 1$ in Lemma 4.17 we get that G contains a set U of t mutual $2k(t-1)$-near-twins. Pick any $v \in U$. Then all $u \in U$ are $2k(t-1)$-near-twins of v, so $|N(v) \triangle N(u)| \leqslant 2k(t-1)$ for all $u \in U$.

We show that $|N(v)| < 2k(t-1)^2 + 2t$. Suppose that $|N(v)| \geqslant 2k(t-1)^2 + 2t$. Then the set $W := \bigcap_{u \in U} N(u)$ has at least $2t$ elements, and so $W - U$ has at least t elements. As every $u \in U$ is adjacent to every $w \in W$ and U and W are disjoint, we have a copy of $K_{t,t}$ as a subgraph of G, a contradiction. Hence, $|N(v)| < 2k(t-1)^2 + 2t < 2kt^2$. □

5.2 Radius-r flip-width and r-admissibility

We have seen that every weakly sparse class $\mathcal{C}$ has bounded degeneracy if and only if $\text{fw}_1(\mathcal{C}) < \infty$. It is known that a weakly sparse class $\mathcal{C}$ has bounded clique-width if and only if $\mathcal{C}$ has bounded tree-width, as clique-width and treewidth are functionally equivalent in weakly sparse classes [GW00]. Hence, for weakly sparse classes $\mathcal{C}$, $\text{fw}_\infty(\mathcal{C}) < \infty$ if and only if $\text{copwidth}_\infty(\mathcal{C}) < \infty$. The general theme is that a dense graph parameter is often equivalent to its sparse counterpart in weakly sparse classes. We now show that for weakly sparse classes, bounded flip-width is indeed equivalent to bounded expansion:

Theorem 5.2. *A class $\mathcal{C}$ has bounded expansion if and only if $\mathcal{C}$ is weakly sparse and has bounded flip-width.*

We prove a more precise result, from which Theorem 5.2 follows.

Theorem 5.3. *There is a constant $c \geqslant 1$ such that for every $r \geqslant 1$ and graph G we have*

$$\text{adm}_r(G) \leqslant (r \cdot \text{fw}_r(G) \cdot \text{degeneracy}(G))^c.$$

We first show how Theorem 5.3 implies Theorem 5.2.

Proof of Theorem 5.2. For the forward direction, assume $\mathcal{C}$ has bounded expansion. Then it has bounded flip-width by Theorem 4.4. Also, as remarked before Theorem 4.4, every class with bounded expansion is weakly sparse.

Conversely, suppose $\mathcal{C}$ is weakly sparse and has bounded flip-width. In particular, $\text{fw}_1(\mathcal{C}) < \infty$, and by Theorem 5.1, $\mathcal{C}$ has degeneracy bounded by some constant d. By Theorem 5.3, for all $r \geqslant 1$ we have that $\text{adm}_r(\mathcal{C}) = (r \cdot \text{fw}_r(\mathcal{C}) \cdot d)^{O(1)} < \infty$. Hence, $\mathcal{C}$ has bounded expansion. □

In the remainder of Section 5.2 we prove Theorem 5.3. Our proof relies on a result of Dvořák [Dvo18], extending a result of Kühn and Osthus [KO04], which we now recall.

An *exact r-subdivision* of a graph G is the graph obtained by replacing every edge of G by a path of length $r+1$. If every edge is replaced by a path of length at most $r+1$, the resulting graph is an *$\leqslant r$-subdivision* of G. For a graph G, let $\tilde{\nabla}^e_r(G)$

denote the maximum average degree of all graphs H whose *exact* r-subdivision is an *induced* subgraph of G.

The following is [Dvo18, Lemma 9].

Lemma 5.4. *For every $r, k, d \geqslant 1$ there is a number $s = s(r, k, d) = O(rdk)^{12}$ such that for every graph G, if $\tilde{\nabla}_0(G) \leqslant d$ and $\tilde{\nabla}^e_r(G) < k$, then $\tilde{\nabla}_r(G) < \tilde{\nabla}_{r-1}(G) + s$.*

An easy induction on r yields the following.

Corollary 5.5. *For every $r, d \geqslant 1$ and d-degenerate graph G,*

$$\tilde{\nabla}_r(G) = O(dr \cdot \sum_{t=1}^{r} \tilde{\nabla}^e_t(G))^{12}.$$

It is well known that every graph of average degree at least d contains a subgraph with minimum degree at least $d/2$. The following proposition yields Theorem 5.3, using Corollary 5.5 and Fact 2.5.

Proposition 5.6 (∗). *Fix $r \geqslant 2, k \geqslant 1$. Let G be the exact $(r-1)$-subdivision of some graph H with minimum degree at least $2rk$. Then $\mathrm{fw}_r(G) > k$.*

Proposition 5.6 is proved in Appendix F, by showing that the vertices of G that correspond to the vertices of H (those of degree larger than 2) form a (r, k, k)-hideout in G. We now show how this proves Theorem 5.3.

Proof of Theorem 5.3. Since $\mathrm{fw}_r(G)$ is monotone in r, by Proposition 5.6, if $\mathrm{fw}_r(G) \leqslant k$ then for all $t < r$, the graph G does not contain an induced t-subdivision of a graph of minimum degree at least $2rk$. Hence, $\tilde{\nabla}^e_t(G) \leqslant 4rk$ for all $t < r$.

If G is d-degenerate then by Corollary 5.5 we have:

$$\tilde{\nabla}_{r-1}(G) = O(dr \sum_{1 \leqslant t < r} \tilde{\nabla}^e_t(G))^{12} \leqslant O(dr^3k)^{12} \leqslant O(drk)^{36}.$$

By Fact 2.5, we get $\mathrm{adm}_r(G) \leqslant r \cdot O(drk)^{108} \leqslant O(drk)^{109}$. □

Proposition 5.6 easily implies the following corollary (see Appendix B.3 for a proof), yielding inequality (6) in Theorem 4.21.

Corollary 5.7 (∗). *If G is the exact 1-subdivision of an n-clique, then $\mathrm{fw}_2(G) > (n-1)/4$. Furthermore, for every graph G, $2\mathrm{VCdim}(G) \leqslant 8\,\mathrm{fw}_2(G) + 2$.*

Hence, weakly sparse classes of bounded flip-width, being exactly the classes with bounded expansion, are by know very well-understood and characterized in multiple ways. For instance, the model checking problem for first-order logic is fixed-parameter tractable for such classes, by the result of Dvořák, Král, and Thomas [DKT13].

6 Flip-width of ordered graphs and twin-width

Let us return to dense graph classes, which are our main focus. As follows from Theorem 4.13, radius-∞ flip-width is functionally equivalent to clique-width, so classes of bounded clique-width are examples of dense graph classes of bounded flip-width. In this section we show that classes of bounded twin-width also have

bounded flip-width, and moreover, twin-width can be characterized in terms of flip-width; however, classes of bounded flip-width strictly extend classes of bounded twin-width. We first recall the definition of twin-width.

An *uncontraction sequence* of a graph G is a sequence $\mathcal{P}_1, \ldots, \mathcal{P}_n$ of partitions of $V(G)$ that starts with the partition $\mathcal{P}_1$ with one part, ends with the partition of $V(G)$ into singletons, and such that every partition $\mathcal{P}_{i+1}$, for $i < n$, is obtained from the previous partition $\mathcal{P}_i$ by splitting one of the parts into two. The *red graph* of a partition $\mathcal{P}$ is the graph whose nodes are the parts of $\mathcal{P}$, and *red edges* connect two parts A, B if A and B are not homogeneous in G. A graph has twin-width at most d if it has an uncontraction sequence such that at every time i, the red graph of the partition $\mathcal{P}_i$ has maximum degree at most d.

First, we prove that $\mathrm{fw}_r(G)$ is bounded in terms of r and the twin-width $\mathrm{tww}(G)$ of G.

Theorem 6.1 ($*$). *Fix $r \in \mathbb{N}$. For every graph G of twin-width d we have:*

$$\mathrm{fw}_r(G) = 2^d \cdot d^{O(r)}.$$

In particular, every class of bounded twin-width has bounded flip-width.

In the proof, the cops use an uncontraction sequence in order to vanquish the robber by constraining him, in round i to some ball of radius r in the red-graph of the partition $\mathcal{P}_i$. In round $i = n$, as every ball of radius r in the red graph of $\mathcal{P}_n$ comprises a single vertex of G, the robber is trapped. We use the bounds from [BFLP23] to get improved bounds. See Appendix G for details.

Therefore, every class of bounded twin-width also has bounded flip-width, but the converse does not hold: the class of subcubic graphs has bounded expansion, and hence bounded flip-width, but does not have bounded twin-width [BGK+21]. So can twin-width be exactly characterized in terms of flip-width?

The twin-width parameter is defined not only for graphs, but also for structures equipped with one or more binary relations. As argued in [BGOdM+22], twin-width may – and perhaps even should – be seen as a parameter of *ordered* graphs, rather than graphs. An ordered graph $G = (V, E, <)$ is equipped with a (symmetric, irreflexive) edge relation E and a total order relation $<$. Every graph can be equipped with some total order without increasing the twin-width, so we may assume the order is present (also such an order can be easily computed from an uncontraction sequence, but the problem of finding it efficiently given the graph G only, remains open).

Similarly as twin-width, the notion of flip-width extends to structures with several binary relations (see Appendix B.4 for details). Theorem 6.1 holds also in the case when G is a binary structure, rather than a graph, with the same proof. In the case of ordered graphs, it is convenient to work with the following variant of flip-width which takes into account that one of the binary relations is a total order. A *k-cut-flip* of an ordered graph $G = (V, E, <)$ is a triple $G' = (V, E', \sim')$, where (V, E') is a (usual) graph that is a k-flip of the graph (V, E), and $\sim'$ is an equivalence relation that partitions V into at most k intervals with respect to $<$. The *weighted graph* associated to G' is the graph with vertices V and edges uv such that $uv \in E'$ or $u \sim' v$, where each edge uv with $u \sim' v$ has weight 0, and the remaining edges have weight 1. Fix a radius $r \in \mathbb{N} \cup \{\infty\}$, a width parameter k, and an ordered graph $G = (V, E, <)$. In the *ordered flip-width game* with radius r

and width k on an ordered graph $G = (V, E, <)$, in round i, the cops announce a k-cut-flip $G_i = (V, E', \sim')$ of the ordered graph $G = (V, E, <)$, and the robber moves from his previous position v_{i-1} to a new position v_i by following a path of total weight at most r in the previous k-cut-flip G_{i-1} (in round $i = 1$, as the cops announce G_1, the robber picks $v_1 \in V$ arbitrarily). The cops win if v_i is isolated in the Gaifman graph of $G_i = (V, E', \sim')$, that is, there is no $w \neq v$ with $vw \in E'$ or $v \sim' w$. The *radius-r ordered flip-width* of an ordered graph $G = (V, E, <)$, denoted $\mathrm{fw}_r^<(G)$, is the smallest number k such that the cops have a winning strategy in the ordered flip-width game with radius r and width k on G.

As mentioned, there is a notion of flip-width for arbitrary structures equipped with binary relations (see Appendix B.4), so in particular, it applies to ordered graphs, seen as relational structures. Write $\mathrm{fw}_r(G)$ for this variant of flip-width of an ordered graph G. The two parameters are functionally equivalent, at the cost of increasing the radius, as stated below, and proved in Appendix G.2.

Lemma 6.2 ($*$). *Fix $r \in \mathbb{N} \cup \{\infty\}$ and an ordered graph $G = (V, E, <)$. Then*

$$\sqrt{\mathrm{fw}_r(G)} \quad \leqslant \quad fw_r^<(G) \quad \leqslant \quad 2\,\mathrm{fw}_{3r+2}(G) + 1.$$

We now prove the main result of Section 6, which says that for ordered graphs $G = (V, E, <)$, the parameters $\mathrm{tww}(G)$ and $\mathrm{fw}_1^<(G)$ are functionally equivalent. In other words, a class of ordered graphs has bounded twin-width if and only if it has bounded radius-one ordered flip-width.

Theorem 6.3. *The following conditions are equivalent for a class $\mathcal{C}$ of ordered graphs:*

1. *$\mathcal{C}$ has bounded twin-width,*
2. *$\mathcal{C}$ has bounded flip-width, that is, $\mathrm{fw}_r(\mathcal{C}) < \infty$ for every $r \in \mathbb{N}$,*
3. *$\mathcal{C}$ has bounded ordered flip-width, that is, $\mathrm{fw}_r^<(\mathcal{C}) < \infty$ for every $r \in \mathbb{N}$,*
4. *$\mathcal{C}$ has bounded radius-one ordered flip-width, that is, $\mathrm{fw}_1^<(\mathcal{C}) < \infty$.*

So for ordered graphs, only the parameters $\mathrm{fw}_\infty^<$ (characterizing bounded clique-width) and $\mathrm{fw}_1^<$ (characterizing bounded twin-width) are relevant, since it follows that $\mathrm{fw}_r^<(G)$ is bounded in terms of $\mathrm{fw}_1^<(G)$ and r, for every $r \in \mathbb{N}$ and ordered graph G.

We obtain the following characterization of twin-width of usual, unordered graphs, in terms of flip-width. Recall that every graph G can be equipped with some total order without increasing the twin-width, and conversely, forgetting a total order of an ordered graph does not increase the twin-width.

Corollary 6.4. *A class $\mathcal{C}$ of graphs has bounded twin-width if and only if every graph in $\mathcal{C}$ can be equipped with a total order, so that the resulting class of ordered graphs has bounded flip-width.*

The implication 1→2 in Theorem 6.1 is by Theorem 6.3 (for binary structures), and the implication 2→3 is by Lemma 6.2. The implication 3→4 is immediate. It remains to prove the implication 4→1, which is done below.

We use a core result of [BGOdM+22], which states that a class $\mathcal{C}$ of ordered graphs has unbounded twin-width if and only if $\mathcal{C}$ has *k-rich divisions*, for every k. A k-rich division of an ordered graph G is a pair of partitions $\mathcal{L}, \mathcal{R}$ of $V(G)$, whose parts are intervals with respect to the order, such that for every interval

$A \in \mathcal{L}$ and k intervals $B_1, \ldots, B_k \in \mathcal{R}$, there are at least k vertices in A with pairwise distinct neighborhoods in $V(G) - (B_1 \cup \cdots \cup B_k)$, and symmetrically, for every interval $B \in \mathcal{R}$ and k intervals $A_1, \ldots, A_k \in \mathcal{L}$, there are at least k vertices in B with pairwise distinct neighborhoods in $V(G) - (A_1 \cup \cdots \cup A_k)$. It is shown in [BGOdM+22, Theorem 21] that if G has no k-rich division, then $\mathrm{tww}(G) = 2^{O(k^2)}$. We now show that a $(k+1)$-rich division can be employed by the robber to evade the cops in the ordered flip-width game of radius 1 and width k. The following lemma immediately yields the implication 4→1 in Theorem 6.3, and finishes its proof.

Lemma 6.5. *Let G be an ordered graph with $\mathrm{fw}_1^<(G) \leqslant k$. Then G does not have a $(k+1)$-rich division $\mathcal{L}, \mathcal{R}$. In particular, $\mathrm{tww}(G) = 2^{O(k^2)}$.*

Proof sketch. Suppose G has a $(k+1)$-rich division $\mathcal{L}, \mathcal{R}$. We show that $\mathrm{fw}_1^<(G) > k$, by describing a winning strategy for the robber in the ordered flip-width game with radius one and width k.

The strategy is as follows: in round i, when the cops announce a new k-cut-flip $G_i = (V, E_i, \sim_i)$ of G, the robber always moves to a vertex v_i in one of the parts of $\mathcal{L}$ (in even-numbered rounds) or of $\mathcal{R}$ (in odd-numbered rounds) that is contained in a single $\sim_i$-equivalence class. We show that the robber can always make such a move, by following a path of weight 1 in the previous k-cut-flip G_{i-1} of G, from his previous position v_{i-1}.

By inductive assumption, suppose that A is contained in a single equivalence class of the previous k-cut-flip G_{i-1}. Since $\sim_i$ partitions $V(G)$ into at most k intervals and $\mathcal{R}$ is a partition into intervals, there are at most $k+1$ intervals $B_1, \ldots, B_{k+1}$ of $\mathcal{R}$ that are not contained in any $\sim_i$-equivalence class. By the condition of a $(k+1)$-rich division there are at least $k+1$ vertices in A with pairwise distinct neighborhoods in $B' := V(G) - (B_1 \cup \cdots \cup B_k)$ in G. This implies in particular that in the previous k-cut-flip G_{i-1} of G, there is an edge joining some vertex $a' \in A$ with some vertex in $b \in B'$. The robber may thus move to b, by a the path $v_{i-1} - a' - b$ of weight 1 in G''. The robber moves to $v_i := b$, maintaining the invariant. □

Reassuming, the duality result of [BGOdM+22] proving the equivalence of unbounded twin-width and having k-rich divisions for all k, can be seen as a min-max theorem for the flip-width game of radius 1, for ordered graphs. Moreover, we now see that degeneracy and twin-width are two flip sides of the same coin: $\mathrm{fw}_1(G)$ corresponds to the degeneracy of G for weakly sparse graphs (see Theorem 3.4 and Theorem 5.2), and $\mathrm{fw}_1^<(G)$ corresponds to the twin-width of G for ordered graphs (see Theorem 6.3). Similarly, classes of bounded flip-width coincide with classes of bounded expansion in the weakly sparse case, and with classes of bounded twin-width in the ordered case. Both of those cases is by now well-understood from an algorithmic, combinatorial, and logical perspective. In particular, the model checking problem for first-order logic is fixed-parameter tractable in each of the two special cases.

7 Closure under transductions

As we have seen, classes of bounded flip-width include all classes of bounded expansion and all classes of bounded twin-width, and characterize those notions in the weakly sparse and totally ordered settings. We argue that classes of bounded flip-width are well behaved, in the sense of having desired closure properties. For instance, if two classes $\mathcal{C}$ and $\mathcal{D}$ have bounded flip-width, then their union $\mathcal{C} \cup \mathcal{D}$ also has bounded flip-width. Other such properties include: closure under disjoint unions (see Example 4.9), and closure under substitution (see Lemma 4.11).

An entire family of closure properties is provided by the notion of first-order *interpretations* or *transductions*. As a very special instance, we saw in Example 4.6 that if $\mathcal{C}$ has bounded flip-width, then the class of edge-complements of graphs from $\mathcal{C}$ also has bounded flip-width. What about, say, the class of squares of graphs from $\mathcal{C}$? (The square of a graph G has vertices $V(G)$ and edges uv such that u and v have a common neighbor in G.) We show that this class also has bounded flip-width, and a similar result holds for every operation that can be defined by a first-order formula, as we now describe. We phrase our result in greater generality for colored graphs.

7.1 Preservation of flip-width under transductions

We start with defining interpretations and transductions, and then state the main result of this section.

Colored graphs Recall that a c-colored graph is a graph together with an assignment of colors from $\{1,\ldots,c\}$ to its vertices. For a c-colored graph G with underlying graph G_0, define $\mathrm{fw}_r(G)$ as $\mathrm{fw}_r(G_0)$ for $r \in \mathbb{N} \cup \{\infty\}$. Say that a class $\mathcal{C}$ of c-colored graphs has *bounded flip-width* if the underlying class of uncolored graphs has bounded flip-width. A c-colored graph is seen as a structure over the signature consisting of the binary relation $E(x,y)$ denoting adjacency, as well as unary predicates $C_1(x),\ldots,C_c(x)$ denoting the respective colors.

Interpretations The following notion is a special case of a (simple, domain-preserving) first-order interpretation. Let G be a c-colored graph and $\varphi(x,y)$ be a first-order formula in the signature of c-colored graphs. Define the graph $\varphi(G)$ with vertices $V(G)$ and edges uv such that $u \neq v$ and $\varphi(u,v) \vee \varphi(v,u)$ holds in G. For a class $\mathcal{C}$ of c-colored graphs, denote $\varphi(\mathcal{C}) := \{\varphi(G) \mid G \in \mathcal{C}\}$. The class $\varphi(\mathcal{C})$ is called an *interpretation* of $\mathcal{C}$, via φ. For example, for the formula $\varphi(x,y) = \neg E(x,y)$ and a graph G, the graph $\varphi(G)$ is the complement $\bar{G}$ of G. And for the formula $\varphi(x,y) = E(x,y) \vee \exists z.[E(x,z) \wedge E(z,y)]$, the graph $\varphi(G)$ is the square of G.

Transductions Say that a class $\mathcal{C}$ *transduces* a class $\mathcal{D}$, or that $\mathcal{D}$ is a *transduction* of $\mathcal{C}$, if there is some $c \geqslant 1$ and class $\widehat{\mathcal{C}}$ of c-colored graphs which is a c-coloring of $\mathcal{C}$, and some first-order formula $\varphi(x,y)$ in the signature of c-colored graphs, such that every graph in $\mathcal{D}$ is an induced subgraph of some graph in $\varphi(\widehat{\mathcal{C}})$ (that is, $\mathcal{D}$ is contained in the hereditary closure of $\varphi(\widehat{\mathcal{C}})$).

Example 7.1. Let $\mathcal{C}$ be the class of all half-graphs (see Figure 1), where the half-graph H_n of order n has vertices $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$, and edges $a_i b_j$ for $1 \leqslant i < j \leqslant n$. We show that $\mathcal{C}$ transduces the class $\mathcal{D}$ consisting of disjoint unions of cliques.

We use two colors. Let $\widehat{\mathcal{C}}$ be the class of all 2-colored half-graphs. Consider the formula $\varphi(x,y)$ expressing that there is no vertex of color 2 which is adjacent to one of x, y, and not the other:

$$\varphi(x,y) \equiv \forall z. C_2(z) \to (E(x,z) \leftrightarrow E(y,z)).$$

We argue that the hereditary closure of $\varphi(\widehat{\mathcal{C}})$ contains $\mathcal{D}$, implying that $\mathcal{D}$ transduces in $\mathcal{C}$.

Let $F \in \mathcal{D}$ be a disjoint union of cliques. Let $1, \ldots, n$ be the vertices of F, for some $n \geqslant 0$, and assume that every connected component of F consists of consecutive vertices in the usual order $1 < \ldots < n$. Consider the half-graph H_n with vertices $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$, and color a vertex b_i with color 2 if i is the largest element of its connected component in F, and with color 1 otherwise. Now, for all $1 \leqslant i, j \leqslant n$, $H_n \models \varphi(a_i, a_j)$ if and only if i and j are adjacent in F. Hence, $\varphi(H_n)[\{a_1, \ldots, a_n\}]$ is isomorphic to F, and therefore F is an induced subgraph of some graph in $\varphi(\widehat{\mathcal{C}})$.

Transductions and flip-width We prove the following theorem.

Theorem 7.2 (∗)**.** *Fix $q \geqslant 0$. There is a computable function $T_q \colon \mathbb{N} \to \mathbb{N}$ with the following property. Fix numbers $r, c \geqslant 1$ and a first-order formula $\varphi(x,y)$ of quantifier rank q in the signature of c-colored graphs. Set $r' := 2^q \cdot r$. Then for every c-colored graph G we have*

$$\mathrm{fw}_r(\varphi(G)) \leqslant T_q(\mathrm{fw}_{r'}(G) \cdot c). \tag{8}$$

In particular, if $\mathcal{C}$ has bounded flip-width, then $\varphi(\mathcal{C})$ has bounded flip-width. Moreover, $T_0(k) = k$ for every k.

The following is an immediate consequence of Theorem 7.2, and the fact that radius-r flip-width is monotone with respect to induced subgraphs.

Corollary 7.3. *If a class $\mathcal{C}$ has bounded flip-width and transduces a class $\mathcal{D}$, then $\mathcal{D}$ has bounded flip-width.*

Since the class of all graphs has unbounded flip-width by Corollary 4.20, we get the following.

Corollary 7.4. *If $\mathcal{C}$ is a class of bounded flip-width, then $\mathcal{C}$ does not transduce the class of all graphs.*

In the language of model theory, Corollary 7.4 says that classes of bounded flip-width are *monadically dependent*, see Section 9. In the phrasing of [GPT22, BNdMS22a], Corollary 7.3 says that classes of bounded flip-width form a *transduction ideal*, whose weakly sparse part, by Theorem 5.2, consists exactly of classes of bounded expansion.

Corollary 7.3 immediately implies that classes with *structurally bounded expansion* [GKN+20], that is, transductions of classes with bounded expansion, have bounded flip-width. This gives another source of examples of dense graph classes of bounded flip-width.

Corollary 7.5. *Every class of structurally bounded expansion has bounded flip-width.*

Corollary 7.3 implies (using Corollary 6.4) the result of [BKTW20], that transductions preserve classes of bounded twin-width. Together with Theorem 5.2, it implies a consequence of [GKN+18], that every weakly sparse class with structurally bounded expansion has bounded expansion.

The proof of Theorem 7.2, which is sketched below, relies on locality of first-order logic, a central notion for analysing first-order formulas on sparse graphs [See96], which, in some form, also plays a key role in understanding first-formulas on classes of bounded twin-width [BKTW20, GPPT22]. The function $T_q(k)$, although computable, is astronomical:

$$T_q(k) := \underbrace{2^{2^{\cdot^{\cdot^{2^m}}}}}_{\text{height } q}$$

where m is the number of distinct k-colored graphs with vertex set $\{1, \ldots, q+1\}$. However, for $q = 0$ we have $T_0(k) = k$, so the bound (8) becomes $\mathrm{fw}_r(\varphi(G)) \leqslant \mathrm{fw}_r(G) = c \cdot \mathrm{fw}_r(G_0)$ in the case when $\varphi(x, y)$ is a quantifier-free formula and G is a c-colored graph with underlying graph G_0.

CMSO **transductions and radius-∞ flip-width** For the case of radius $r = \infty$, we prove that fw_∞ is preserved under transductions expressed in the more powerful logic called CMSO. *Counting Monadic Second Order Logic* (CMSO) is the extension of first-order logic, where apart from first-order quantifiers $\exists x, \forall x$, that range over vertices x of a graph, we have second-order quantifers $\exists X, \forall X$, that range over sets X of vertices of the graph. We additionally have the atomic predicate $x \in X$ that allows to check whether a given vertex belongs to a given set, and the divisibility predicates $\mathrm{div}_k(X)$, for $k \geqslant 1$, where $\mathrm{div}_k(X)$ holds for a given set of vertices X if and only if $|X|$ is divisible by k. The usual constructs of first-order logic (boolean connectives and relation symbols, such as adjacency in a graph) are also included.

CMSO is able to express non-local properties. For instance, the formula

$$\varphi(x, y) = \neg\exists X. \left[(x \in X) \land (y \notin Y) \land \forall z. \forall t. \left[(z \in X) \land E(z, t) \to (t \in X)\right]\right]$$

expresses that x and y lie in the same connected component.

Still, CMSO enjoys property similar to locality of first-order logic, called compositionality, which is an analogue of locality for the radius $r = \infty$. With the same proof as in Theorem 7.2, we get the following.

Theorem 7.6 ($*$)**.** *Let $\mathcal{C}$ be a class of c-colored graphs of bounded ∞-flip-width and let $\varphi(x, y)$ be a formula of* CMSO*. Then $\varphi(\mathcal{C})$ has bounded ∞-flip-width.*

Say that a class $\mathcal{C}$ CMSO-*transduces* a class $\mathcal{D}$, or that $\mathcal{D}$ is a CMSO-*transduction* of $\mathcal{C}$, if there is some $c \geqslant 1$ and class of c-colored graphs $\widehat{\mathcal{C}}$ which is a c-coloring of $\mathcal{C}$, and some CMSO formula $\varphi(x, y)$ in the signature of c-colored graphs, such that every graph in $\mathcal{D}$ is an induced subgraph of some graph in $\varphi(\widehat{\mathcal{C}})$ (that is, $\mathcal{D}$ is contained in the hereditary closure of $\varphi(\widehat{\mathcal{C}})$).

Corollary 7.7. *If $\mathcal{C}$ is a class with $\mathrm{fw}_\infty(\mathcal{C}) < \infty$ that* CMSO*-transduces a class $\mathcal{D}$, then $\mathrm{fw}_\infty(\mathcal{D}) < \infty$.*

As shown in Section 4.2, together with the result of Courcelle and Oum, this proves Theorem 4.13, that classes of bounded clique-width are exactly those with bounded fw_∞.

Before giving the details about the proof of Theorem 7.2, we formulate a notion that will be also useful in other contexts.

7.2 Transferring strategies

The following notion allows to transfer a winning strategy of the cops from a graph G to a graph H. Let G and H be two graphs with $V(H) \subseteq V(G)$. Fix $r_G, r_H \in \mathbb{N} \cup \{\infty\}$ and $k, \ell \in \mathbb{N}$. Suppose furthermore that the following are given:

- a mapping F that maps each k-flip G' of G to an ℓ-flip $H' = F(G')$ of H,
- a strategy of the cops in the game on the graph G with radius r' and width k.

This induces the following strategy of the cops in the flip-width game of radius r_H and width ℓ on the graph H. Consider a play of this game, and call it the *H-game*. Simultaneously, initiate the *G-game* on G, with radius r_G and width k, in which we will copy robber's moves from the H-game. Whenever the cops announce a k-flip G' of G in the G-game, then in the H-game the cops announce the ℓ-flip $H' = F(G')$ of H. Whenever the robber moves to a vertex v in the H-game, we also move the robber to v in the H-game. Note that this move might not be a valid move in the H-game. In this case, the cops in the G-game cannot proceed further with copying their moves from the H-game, so the cops announce their defeat in the G-game. We say the resulting strategy of the cops (in the G-game) is *transferred* from the considered strategy in the H-game, according to the mapping F.

The following lemma gives a condition which implies that if the original strategy on G is winning, then the transferred strategy on H is winning.

Lemma 7.8. *Fix $r \in \mathbb{N} \cup \{\infty\}$ and $k, \ell, s \geqslant 1$. Let H, G be two graphs with $V(H) \subseteq V(G)$. Suppose that for all $k \geqslant 1$ and every k-flip G' of G there is some ℓ-flip $H' = F(G')$ of H, such that*

$$\mathrm{dist}_{G'}(u,v) \leqslant s \qquad \text{for all } uv \in E(H'). \tag{9}$$

Then, transferring a winning strategy of the cops from the flip-width game on G with radius rs and width k, according to the mapping F, results in a winning strategy of the cops in the flip-width game on H with radius r and width ℓ. In particular,

$$\mathrm{fw}_{rs}(G) \leqslant k \qquad \textit{implies} \qquad \mathrm{fw}_r(H) \leqslant \ell.$$

In the case $r = \infty$, the conclusion reads "$\mathrm{fw}_\infty(G) \leqslant k$ implies $\mathrm{fw}_\infty(H) \leqslant \ell$". Moreover we can then replace $\leqslant s$ by $< \infty$ in (9), as the statement then does not depend on s, and setting $s := |V(G)|$ and $\mathrm{dist}_{G'}(u,v) \leqslant s$ implies $\mathrm{dist}_{G'}(u,v) < \infty$.

Proof. Consider a play in the flip-width game on H of radius r and width ℓ, according to the strategy transferred from G, as described above. Suppose that, in some round, G' is the announced flip in the G-game and $H' = F(G')$ is the announced flip in the H-game, and that v is the new vertex chosen by the robber in

the H-game. By (9), the following holds:

$$B^r_{H'}(v) \subseteq B^{rs}_{G'}(v).$$

In the case $r = \infty$, each side of the inclusion should be interpreted as the connected component of v in the appropriate graph.

In particular, in the next round, every valid move of robber in the H-game will also be a valid move of robber in the G-game (this is trivially satisfied in the first round), and if robber is trapped in the G-game, that is, $|B^{rs}_{G'}(v)| = 1$, then also $|B^r_{H'}(v)| = 1$, so he is also trapped in the H-game. Thus, this describes a winning strategy for the cops in the flip-width game on H with radius r and width ℓ. In particular, $\mathrm{fw}_r(H) \leqslant \ell$. □

7.3 Proof of Theorem 7.2

We now sketch the proof of Theorem 7.2. The details are presented in Appendix H, along with the proof of Theorem 7.6.

For simplicity, we assume the case $c = 1$, that is, when the considered graphs G have no colors. The general case proceeds analogously.

Our aim is to transfer a winning strategy of the cops from G to $H = \varphi(G)$ (with appropriate radii), by applying Lemma 7.8. So for every flip G' of G we need to produce a flip $\varphi(G)'$ of $\varphi(G)$ such that adjacent vertices in $\varphi(G)'$ are not too far in G'. We first show how to achieve this in the case when $G' = G$, using locality, a well known tool from finite model theory, which we now recall.

Locality Fix a number $s \in \mathbb{N}$. Say that a formula $\varphi(x,y)$ is *s-local*, if for any graph G there is a labelling of the vertices of G using a bounded number of labels (depending only on φ, and not on G) such that for any two vertices u, v of G with $\mathrm{dist}_G(u,v) > s$, whether or not $G \models \varphi(u,v)$ depends only on the label of u and the label of v. It is well-known (and follows for instance from Gaifman's locality theorem) that every formula $\varphi(x,y)$ of first-order logic is s-local for some radius s depending only on φ. Namely, one can take $s = 2^q$, where q is the quantifier rank of φ. The label assigned to a vertex v of G as above, is essentially the set of formulas $\alpha(x)$ of quantifier rank at most q, such that $\alpha(v)$ holds in G. The number of such formulas, up to equivalence, is finite, and is bounded by $T_q(1)$, where T_q is the function described above.

Flipping $\varphi(G)$ We now show how to obtain a flip $\varphi(G)'$ of $\varphi(G)$ such that vertices that are adjacent in $\varphi(G)'$ are not too far in G. Let $\mathcal{P}$ be the partition of $V(G) = V(\varphi(G))$ such that two vertices of G are in the same part if they get the same label. In particular, $|\mathcal{P}| \leqslant T_q(1)$. Now, in the graph $\varphi(G)$ flip a pair of parts A, B of $\mathcal{P}$ if and only if there is a pair of vertices $u \in A$ and $v \in B$, such that $\mathrm{dist}_G(u,v) > s$ and $G \models \varphi(u,v)$, equivalently, $uv \in E(\varphi(G))$. The statement above implies that whether or not we flip A and B, does not depend on the choice of $u \in A$ and $v \in B$ such that $\mathrm{dist}_G(u,v) > s$. This yields a $\mathcal{P}$-flip of $\varphi(G)$, which we denote $\varphi(G)'$. Then the following holds for all $u, v \in V(G)$:

$$uv \in E(\varphi(G)') \qquad \text{implies} \qquad \mathrm{dist}_G(u,v) \leqslant s.$$

We generalize the above reasoning to the case when G' is a k-flip of G. Again, the goal is to construct a flip $\varphi(G)'$ of $\varphi(G)$ such that vertices that are adjacent in $\varphi(G)'$ are not too far in G'.

We treat G' as a k-colored graph, by adding colors that mark parts of the partition that is used to produce the k-flip G' of G. The key observation is that we can write a formula $\varphi'(x,y)$ that makes use of those colors, and such that $G' \models \varphi'(u,v)$ if and only if $G \models \varphi(u,v)$, for any pair of vertices $u,v \in V(G)$. This is because we can write a formula $\varepsilon(x,y)$ such that $G' \models \varepsilon(u,v)$ if and only if $G \models E(u,v)$ (the formula $\varepsilon(x,y)$ checks the colors of x and y and whether x and y are adjacent). The formula $\varphi'(u,v)$ is obtained by replacing each atomic formula $E(z,t)$ with $\varepsilon(z,t)$. Applying the same argumentation as above, but this time to the formula φ' and the k-colored graph G', we obtain a $\mathcal{P}$-flip $\varphi(G)'$ such that the following holds for all $u,v \in V(G)$:

$$uv \in E(\varphi(G)') \qquad \text{implies} \qquad \mathrm{dist}_{G'}(u,v) \leqslant s.$$

Moreover, the size of $\mathcal{P}$ can be bounded by a number ℓ depending on the formula φ and the number k. Since this holds for every k-flip G' of G, we can now apply Lemma 7.8 and conclude that $\mathrm{fw}_r(\varphi(G)) \leqslant \ell$.

8 Definable flip-width

Determining the flip-width of radius r of a given graph G seems computationally difficult. The space of all configurations in the flip-width game of radius r and width k has size exponential in $|G|$, and the naive algorithm for determining whether $\mathrm{fw}_r(G) \leqslant k$, which explores the space of all configurations, therefore runs in time exponential in $|G|$. We expect that for possible algorithmic applications, an algorithm which approximates $\mathrm{fw}_r(G)$, instead of computing it exactly, should suffice. This is what happens in the case of weak coloring numbers, and of twin-width, where such an approximation algorithm is not known in general, but is known in various special cases.

In this section, we consider the definable flip-width game of radius r and width k, in which the moves allowed for the cops are parameterized by tuples of vertices of the graph, rather than by partitions of the vertex set. Among other things, this reduces the computational complexity of determining the flip-width of a given graph. The main result of this section is Theorem 8.4 that says that the definable flip-width at radius r can be bounded in terms of the flip-width at radius $5r$. This allows to obtain, in Theorem 8.7, an algorithm for approximating flip-width, which runs in time $n^{O(k)} \cdot O_k(1)$, where k is the flip-width. We use tools related to VC-dimension to accomplish this.

Atomic types and definable flips Let S be a set of vertices of a graph G. Consider the partition $\mathcal{P}$ of $V(G)$ such that two vertices $u,v \in V(G)$ are in the same part of $\mathcal{P}$ if either $u = v = s$ for some $s \in S$, or $u,v \notin S$ and $N(u) \cap S = N(v) \cap S$. The equivalence classes of the partition $\mathcal{P}$ are called *S-types*.

Say that a graph G' is an *S-definable flip of G* if G' is a $\mathcal{P}$-flip of G, where $\mathcal{P}$ is the partition of $V(G)$ into S-types. Say that G' is a *k-definable flip* of G, if G' is an S-definable flip of G for some $S \subseteq V(G)$ with $|S| \leqslant k$.

Note that a k-definable flip of G is a $(2^k + k)$-flip of G. However, there is no function f such every k-flip of a graph G is a $f(k)$-definable flip of G. For instance, the graph G_n obtained from an n-clique K_n by adding n isolated vertices, is a 2-flip of K_{2n}, but is not a k-definable flip of K_{2n} for $k < n$.

Definable flip-width game Fix $r \in \mathbb{N} \cup \infty$. The *definable flip-width game of radius r and width k* is defined in the same way as the flip-width game of radius r, but now in each round the cops are allowed to announce a k-definable flip G' of G, rather than a k-flip of G.

Definition 8.1. *Fix $r \in \mathbb{N} \cup \infty$. The* radius-$r$ definable flip-width *of a graph G, denoted* $\mathrm{dfw}_r(G)$*, is the smallest number k such that the cops have a winning strategy in the definable flip-width game of radius r and width k on G.*

As every k-definable flip of G is a $(2^k + k)$-flip of G, it follows that for every $r \in \mathbb{N} \cup \{\infty\}$ and graph G we have:

$$\mathrm{fw}_r(G) \leqslant 2^{\mathrm{dfw}_r(G)} + \mathrm{dfw}_r(G). \tag{10}$$

One advantage of the definable version of the flip-width game is that it has far fewer configurations than the original flip-width game. As there are only $O(n^{k+1} \cdot 2^{4^k})$ configurations in the definable flip-width game of width k, we get that it can be decided in time $n^{O(k)} \cdot 2^{O(2^k)}$ whether a given n-vertex graph has $\mathrm{dfw}_r(G) \leqslant k$ (see Appendix B.5 for a proof).

Lemma 8.2 (∗). *There is an algorithm that, given a graph G and numbers $k \in \mathbb{N}$ and $r \in \mathbb{N} \cup \{\infty\}$, determines whether* $\mathrm{dfw}_r(G) \leqslant k$ *in time* $n^{O(k)} \cdot 2^{O(2^k)}$.

The following lemma is an immediate consequence of the Sauer-Shelah-Perles lemma (Lemma 2.8).

Lemma 8.3 (∗). *Fix $r \in \mathbb{N} \cup \{\infty\}$. For every graph G we have:*

$$\mathrm{fw}_r(G) = O(\mathrm{dfw}_r(G)^{\mathrm{VCdim}(G)}). \tag{11}$$

The following is the main result of Section 8. It says that the definable flip-width can be bounded in terms of the flip-width, at the cost of increasing the radius. Below, $5 \cdot \infty$ is interpreted as ∞.

Theorem 8.4. *Fix $r \in \mathbb{N} \cup \{\infty\}$. For every graph G we have:*

$$\mathrm{dfw}_r(G) = O(\mathrm{fw}_{5r}(G)^3). \tag{12}$$

We prove Theorem 8.4 below. We first observe some consequences. The bounds (10) and (12) give the following.

Corollary 8.5. *The following conditions are equivalent for a graph class $\mathcal{C}$:*

1. *$\mathcal{C}$ has bounded flip-width, that is,* $\mathrm{fw}_r(\mathcal{C}) < \infty$ *for all $r \geqslant 1$,*
2. $\mathrm{dfw}_r(\mathcal{C}) < \infty$ *for all $r \geqslant 1$.*

Similarly, for the case $r = \infty$, using Theorem 4.13 we get the following characterization of classes of bounded clique-width in terms of the definable flip-width game with radius ∞:

Corollary 8.6. *The following conditions are equivalent for a graph class $\mathcal{C}$:*

1. *$\mathcal{C}$ has bounded clique-width,*
2. $\mathrm{dfw}_\infty(\mathcal{C}) < \infty$.

Finally, we get an algorithm for approximating the flip-width of a given graph. The algorithm is an approximation algorithm, that is, unlike the algorithm in Lemma 8.2, it does not allow to exactly determine whether the radius-r flip-width of a given graph G is smaller than a given number k. Rather, it recognises one of two, non-exclusive, cases: whether $\mathrm{fw}_r(G)$ is small comparing to k, and whether $\mathrm{fw}_{5r}(G)$ is large comparing to k. Note that there is a gap in the radii, r and $5r$. The running time of the algorithm is $O_k(1) \cdot n^{O(k)}$, which is called an XP algorithm (parameterized by k) in the language of parameterized complexity.

Theorem 8.7. *There is a constant $C > 0$ and an algorithm that inputs a graph G and numbers $r, k \in \mathbb{N}$, runs in time $n^{O(k)} \cdot 2^{O(2^k)}$, and either concludes that $\mathrm{fw}_r(G) \leqslant 2^k + k$, or concludes that $\mathrm{fw}_{5r}(G) \geqslant C \cdot k^{1/3}$.*

Proof. The algorithm tests whether $\mathrm{dfw}_r(G) \leqslant k$ in time $n^{O(k)} \cdot 2^{O(2^k)}$, using Lemma 8.2. If $\mathrm{dfw}_r(G) \leqslant k$ it concludes that $\mathrm{fw}_r(G) \leqslant 2^k + k$, by (10). If $\mathrm{dfw}_r(G) > k$, it concludes that $\mathrm{fw}_{5r}(G) \geqslant Ck^{1/3}$ by (12), where $C > 0$ is some fixed constant. □

Proof of Theorem 8.4 We now turn to the proof of Theorem 8.4. We use a result concerning graphs of small VC-dimension. Recall that $\mathrm{VCdim}(G) = O(\mathrm{fw}_1(G))$ by Theorem 4.21. The following result from [BDG+22], relies on the (p, q)-theorem of Matoušek [Mat04].

Lemma 8.8. *Fix $k, d \in \mathbb{N}$. Let V be a set equipped with:*

- *a binary relation $E \subseteq V \times V$ of VC-dimension at most d,*
- *a pseudometric $\mathrm{dist}\colon V \times V \to \mathbb{R}_{\geqslant 0} \cup \{\infty\}$,*
- *and a partition $\mathcal{P}$ of size at most k,*

such that $E(u, v)$ depends only on the $\mathcal{P}$-class of u and the $\mathcal{P}$-class of v whenever $\mathrm{dist}(u, v) > 1$. Then there is a set $S \subseteq V$ of size $O(dk^2)$, such that $E(u, v)$ depends only on the S-types of u and of v, whenever $\mathrm{dist}(u, v) > 5$.

We reformulate Lemma 8.8 in terms of flips, as follows.

Corollary 8.9. *Let G be a graph and $d = \mathrm{VCdim}(G)$, and let G' be a k-flip of G. Then there is a $O(dk^2)$-definable flip G'' of G such that*

$$\mathrm{dist}_{G'}(u, v) \leqslant 5 \qquad \text{for all } u, v \in E(G'') \tag{13}$$

Proof. Apply Lemma 8.8 to $V = V(G)$, $E = E(G)$, $\mathrm{dist}\colon V \times V \to \mathbb{N} \cup \{\infty\}$ denoting the shortest path metric in G', and the partition $\mathcal{P}$ with $|\mathcal{P}| \leqslant k$ such that G' is a $\mathcal{P}$-flip of G. Note that for any two vertices $u, v \in V$, whether or not $uv \in E(G)$ holds, can be determined basing only on the $\mathcal{P}$-class of u, the $\mathcal{P}$-class of v, and on the information whether $uv \in E(G')$ holds. In particular, $uv \in E(G)$ depends only on the $\mathcal{P}$-class of u and the $\mathcal{P}$-class of v, for all $u, v \in V$ that are not adjacent in G', equivalently, with $\mathrm{dist}(u, v) > 1$. Hence, the assumption of Lemma 8.8 is satisfied.

Let S with $|S| = O(dk^2)$ be as in the conclusion of the lemma, so that whether or not $uv \in E(G)$, depends only on the S-type of u and the S-type of v, for all $u, v \in V$ with $\mathrm{dist}(u, v) > 5$.

Let $\mathcal{S}$ denote the partition of $V(G)$ into S-types. Let G'' be the $\mathcal{S}$-flip of G that flips between two S-types A and B if and only if there are some $u \in A, v \in B$ with $\mathrm{dist}(u, v) > 5$ and $uv \in E(G)$. The conclusion follows. □

Theorem 8.4 easily follows from Corollary 8.9.

Proof of Theorem 8.4. Let G be a graph and let $k = \mathrm{fw}_{5r}(G)$. In particular, by Theorem 4.21, we have that $\mathrm{VCdim}(G) = O(k)$. By Corollary 8.9, the assumptions of Lemma 7.8 (with $H := G$) are satisfied, where instead of an ℓ-flip G' of G, we have a $O(k^3)$-definable flip G' of G (as provided by Corollary 8.9). By transferring the winning strategy of the cops in the flip-width game of radius $5r$ and width k on G, according to the mapping $G' \mapsto G''$ as given by Corollary 8.9, we get a strategy for the cops in flip-width game of radius r on G, which uses only $O(dk^2)$-definable flips. By Corollary 8.9, the assumptions of Lemma 7.8 are satisfied, so this yields a winning strategy of the cops that will use only $O(k^3)$-definable flips. Hence, $\mathrm{dfw}_r(G) = O(k^3)$. □

9 Almost bounded flip-width

Recall from Fact 2.7 that a hereditary graph class $\mathcal{C}$ is nowhere dense if and only if for every $r \geqslant 1$ and $G \in \mathcal{C}$ we have $\mathrm{wcol}_r(G) = |G|^{o(1)}$. Inspired by this characterization, we extend the notion of bounded flip-width as follows.

Definition 9.1. *A graph class $\mathcal{C}$ has* almost bounded flip-width *if for every $r \geqslant 1$ and real $\varepsilon > 0$ we have $\mathrm{fw}_r(G) = O_{\varepsilon,r}(|G|^{\varepsilon})$ for every graph G in the hereditary closure of $\mathcal{C}$.*

Note that we consider all graphs G from the hereditary closure of $\mathcal{C}$. Otherwise, the class consisting of every graph G with $2^{|G|}$ isolated vertices added to it, would have almost bounded flip-width, while according to the above definition, it does not. Indeed, we have the following lemma, which is an immediate consequence of Lemma 4.23.

Lemma 9.2. *Every graph class with almost bounded flip-width has bounded VC-dimension.*

Clearly, every class with bounded flip-width has almost bounded flip-width. As we conjecture (see Conjecture 9.7), classes of almost bounded flip-width coincide with monadically dependent classes (see definition below), analogously to the characterization of nowhere dense classes in Fact 2.7.

In this section, we provide some evidence towards this conjecture. In Theorem 9.8, we prove that a weakly sparse class has almost bounded flip-width if and only if it is nowhere dense, if and only if it is monadically dependent. In Theorem 9.11 we prove that structurally nowhere dense classes have almost bounded flip-width. In Theorem 9.14 we prove that edge-stable classes of almost bounded flip-width are monadically dependent. In Theorem 9.16 we prove that classes of ordered graphs of almost bounded flip-width coincide with classes of bounded twin-width, and with classes of bounded flip-width. We start with recalling the discussed notions.

9.1 Monadic dependence and monadic stability

The following notion, due to Shelah [She86] (see also [BL21]), originates in model theory.

Definition 9.3. *A graph class $\mathcal{C}$ is* monadically dependent *(or monadically NIP) if and only if it does not transduce the class of all graphs.*

Monadically dependent classes have recently attracted attention in areas of structural and algorithmic graph theory [AA14, NdMP⁺21, BGOdM⁺22, GPT22], as it is conjectured (see Conjecture 1.1) that monadically dependent classes are precisely those, for which model checking first-order logic is fixed-parameter tractable.

Monadically dependent classes include all nowhere dense classes, and in fact, among weakly sparse classes, they provide an exact characterization:

Fact 9.4 (Consequence of [Dvo18] and [AA14])**.** *Let $\mathcal{C}$ be a weakly sparse graph class. Then $\mathcal{C}$ is nowhere dense if and only if $\mathcal{C}$ is monadically dependent.*

Monadically dependent classes also include all classes of bounded twin-width [BKTW20], classes of bounded expansion, and more generally, by Corollary 7.4, all classes of bounded flip-width.

An important subfamily of monadically dependent classes consists of monadically stable classes. A class $\mathcal{C}$ is *monadically stable* if it does not transduce the class of all half-graphs (see Fig. 1). A class $\mathcal{C}$ is *edge-stable* if it excludes some half-graph as a semi-induced bipartite graph. More precisely, there is $k \geqslant 1$ such that there do not exist $G \in \mathcal{C}$ and vertices $a_1, \ldots, a_k$ and $b_1, \ldots, b_k$ such that $a_i b_j \in E(G) \iff i < j$ for all $i, j \in \{1, \ldots, k\}$.

The following result is proved in [NdMP⁺21, Theorem 1.3], see also [BL22, Theorem 3.20].

Fact 9.5. *Let $\mathcal{C}$ be an edge-stable graph class. Then $\mathcal{C}$ is monadically stable if and only if it is monadically dependent.*

Monadically stable classes include all nowhere dense classes [AA14], as well as transductions of nowhere dense classes, called *structurally nowhere dense* classes. It is not known whether all monadically stable classes are structurally nowhere dense (this has been conjectured in [NdMP⁺21, Conjecture 6.1]). The class of half-graphs is clearly not monadically stable, and has bounded (linear) clique-width. Hence, monadically stable classes are incomparable with classes of bounded clique-width. They are also incomparable with classes of bounded flip-width, as witnessed by the class of half-graphs on one side, and any nowhere dense class which does not have bounded expansion on the other side.

Monadic dependence can be defined not only for graph classes, but for arbitrary classes of structures, e.g. classes of ordered graphs. A class of ordered graphs is monadically dependent if it does not transduce the class of all graphs, where now the transduction may involve the edge relation symbol, as well as the total order $<$ (and the color predicates). The following result is proved in [BGOdM⁺22].

Fact 9.6. *Let $\mathcal{C}$ be a class of ordered graphs. Then $\mathcal{C}$ is monadically dependent if and only if $\mathcal{C}$ has bounded twin-width.*

As every graph can be equipped with a total order without increasing its twin-width, it follows that a class of (unordered) graphs $\mathcal{C}$ has bounded twin-width if and only if it can be obtained from a monadically dependent class $\mathcal{C}^<$ of ordered graphs, by forgetting the order. In particular, even though the class $\mathcal{C}$ of cubic graphs is monadically dependent (as it is nowhere dense), it cannot be expanded to a monadically dependent class $\mathcal{C}^<$ of ordered graphs, since $\mathcal{C}$ has unbounded twin-width by [BGK+21].

We conjecture that classes of almost bounded flip-width coincides are exactly the monadically dependent classes.

Conjecture 9.7. *A hereditary graph class has almost bounded flip-width if and only if it is monadically dependent.*

Currently, we are unable to prove neither implication in this conjecture. However, In the rest of Section 9, we provide evidence towards this conjecture, by confirming it in restricted settings.

9.2 Weakly sparse classes of almost bounded flip-width

Analogously to Theorem 5.2, which characterizes classes with bounded expansion as exactly the weakly sparse classes of bounded flip-width, we get a characterization of nowhere dense classes in terms of almost bounded flip-width.

Theorem 9.8. *Let $\mathcal{C}$ be a weakly sparse graph class. Then the following conditions are equivalent:*

1. *$\mathcal{C}$ is nowhere dense,*
2. *$\mathcal{C}$ has almost bounded flip-width,*
3. *$\mathcal{C}$ is monadically dependent.*

The equivalence of the first and last condition is by Fact 9.4, so we only prove the equivalence of the first two.

Proof. We first show that every nowhere dense class has almost bounded flip-width. Every nowhere dense class $\mathcal{C}$ is weakly sparse, so excludes some $K_{t,t}$ as a subgraph. By Theorem 4.4, Theorem 3.5, and the forward implication in Fact 2.7, we have that $\mathrm{fw}_r(G) = O_{r,\varepsilon}(n^{t\varepsilon})$ for every n-vertex graph $G \in \mathcal{C}$. Since this holds for every $\varepsilon > 0$ and t is fixed, the conclusion follows by rescaling ε.

Conversely, suppose that $\mathcal{C}$ has almost bounded flip-width, and excludes $K_{t,t}$ as a subgraph. Without loss of generality, we may assume that $\mathcal{C}$ is hereditary. We have $\mathrm{fw}_1(G) = O_\varepsilon(n^\varepsilon)$ for every n-vertex graph $G \in \mathcal{C}$, and by Theorem 5.1, $\mathrm{degeneracy}(G) = O_\varepsilon(n^\varepsilon \cdot 2t^2) = O_\varepsilon(n^\varepsilon)$. By Theorem 5.3, for every $r \geqslant 1$ we have $\mathrm{adm}_r(G) \leqslant (r \cdot \mathrm{fw}_r(G) \cdot \mathrm{degeneracy}(G))^{O(1)} \leqslant (r \cdot O_{r,\varepsilon}(n^\varepsilon) \cdot O_\varepsilon(n^\varepsilon))^{O(1)} \leqslant O_{r,\varepsilon}(n^{O(\varepsilon)})$ for every $\varepsilon > 0$ and n-vertex graph G. By (1), we have $\mathrm{wcol}_r(G) = O_{r,\varepsilon}(n^{O(\varepsilon)})$ for every $r \geqslant 1$ and n-vertex graph $G \in \mathcal{C}$. Therefore, $\mathcal{C}$ is nowhere dense, by the backwards implication in Fact 2.7. $\square$

As there exist nowhere dense classes of unbounded degeneracy, we get the following.

Corollary 9.9. *There is a class that has almost bounded flip-width, but does not have bounded flip-width.*

By taking the substitution closure (see Section 4.1) of the class from the corollary above, we obtain a class which has almost bounded flip-width (by Lemma 4.11), but is not nowhere dense (not even monadically stable), and has unbounded flip-width.

9.3 Structurally nowhere dense classes

We are unable to determine whether classes of almost bounded flip-width are closed under transductions. Observe that the bound $\mathrm{fw}_r(\varphi(G)) \leqslant T_q(\mathrm{fw}_{r'}(G))$ in Theorem 7.2 is not polynomial in $\mathrm{fw}_{r'}(G)$. It is, however, linear in the case when φ is a quantifier-free formulas $\varphi(x,y)$, so we get the following.

Corollary 9.10. *Let $\mathcal{C}$ be a class of k-colored colored graphs of almost bounded flip-width, and let $\varphi(x,y)$ be a quantifier-free formula. Then the class $\varphi(\mathcal{C})$ has almost bounded flip-width.*

Even though we do not know whether classes of almost bounded flip-width are closed under transductions, we confirm that structurally nowhere dense classes have almost bounded flip-width.

Theorem 9.11. *Every structurally nowhere dense class has almost bounded flip-width.*

To prove Theorem 9.11, we use the main result of [DGK+22a], which essentially implies that for every structurally nowhere dense class $\mathcal{C}$ there is an almost nowhere dense class $\mathcal{B}$ of structures equipped with functions, and a quantifier-free formula $\varphi(x,y)$ involving function symbols, such that $\mathcal{C} \subseteq \varphi(\mathcal{B})$. This is made precise below.

Say that a class $\mathcal{C}$ of graphs is *almost nowhere dense* if for every $r \geqslant 1$ and $\varepsilon > 0$ we have $\mathrm{wcol}_r(G) = O_{r,\varepsilon}(|G|^\varepsilon)$, for all $G \in \mathcal{C}$. Crucially, $\mathcal{C}$ does not need to be hereditary, otherwise this notion would coincide with nowhere denseness by Fact 2.7.

Fix a signature Σ consisting of unary relation symbols, binary relation symbols, and unary function symbols. The VC-dimension of a Σ-structure B, denoted $\mathrm{VCdim}(B)$, is the maximum of the VC-dimensions of the binary relations of B (see Section 2.3). Here, the functions of B are ignored.

The following result, apart from the 'moreover' part, is a straightforward consequence of [DGK+22a, Theorem 3]. It says that every structurally nowhere dense class $\mathcal{C}$ interprets in an almost nowhere dense class $\mathcal{B}$ of structures, via a quantifier-free interpretation using a unary function symbol. Moreover, every $G \in \mathcal{C}$ interprets in some $B \in \mathcal{B}$ with $|B| \leqslant O(|G|)$. Finally, the binary relations of the binary structures in $\mathcal{B}$ have bounded VC-dimension, which will be important in the next lemma, for controlling the bounds on the flip-width of G.

Theorem 9.12 ($*$)**.** *Let $\mathcal{C}$ be a structurally nowhere dense graph class. There is a signature Σ consisting of unary and binary relation symbols and one function symbol, a class $\mathcal{B}$ of Σ-structures which is almost nowhere dense, and a quantifier-free symmetric formula $\varphi(x,y)$ with the following property. For every graph $G \in \mathcal{C}$ there is some $B \in \mathcal{B}$ with $|B| = O(|G|)$, such that G is an induced subgraph of $\varphi(B)$. Moreover,* $\mathrm{VCdim}(B) < \infty$.

The 'moreover' part is shown by analysing the construction from [DGK+22a] (see [DGK+22b] for the full version). The following lemma is proved by extending the ideas used in the proof of Theorem 7.2.

Lemma 9.13 (∗)**.** *Let Σ be a signature consisting of unary and binary relation symbols, and unary function symbols. Fix $k, r \geqslant 0$, and a symmetric quantifier-free Σ-formula $\varphi(x, y)$. There are numbers $p = O_\varphi(k)$ and $r' = O_\varphi(r)$ such that the following holds. Let B be a Σ-structure of VC-dimension at most k and G_B be its Gaifman graph. Then*

$$\mathrm{fw}_r(\varphi(B)) = O(\mathrm{copwidth}_{r'}(G_B))^p.$$

The key insight is that a bound d on the VC-dimension implies that if the cops in the cop-width game occupy a set S of vertices of a graph G, then the partition of $V(G)$ into S-types has size $O(|S|^d)$, and this partition is used by the cop in the flip-width game.

The two statements above are proved in Appendix I. Theorem 9.11 follows, as we now show.

Proof of Theorem 9.11. Let $\mathcal{C}$ be a structurally nowhere dense class. Without loss of generality, $\mathcal{C}$ is hereditary. Let $\mathcal{B}$ and $\varphi(x, y)$ be as in Theorem 9.12, and $k \in \mathbb{N}$ be such that $\mathrm{VCdim}(B) < k$ for $B \in \mathcal{B}$.

Fix $\varepsilon > 0$. Let $G \in \mathcal{C}$ and $B \in \mathcal{B}$ be such that G is an induced subgraph of $\varphi(B)$ and $|B| = O(|G|)$, and let G_B be the Gaifman graph of B. Let p and r' be as in Lemma 9.13. Then we have:

$$\mathrm{fw}_r(G) \leqslant \mathrm{fw}_r(\varphi(B)) = O(\mathrm{copwidth}_{r'}(G_B))^p = O_{r',\varepsilon}(|B|^{\varepsilon p}) = O_{r',\varepsilon}(|G|^{\varepsilon p}).$$

Since r' depends only on r and φ, and p is a constant, and $\varepsilon > 0$ is arbitrary, this proves that $\mathcal{C}$ has almost bounded flip-width. □

9.4 Edge-stable classes of almost bounded flip-width

As mentioned, we conjecture (see Conjecture 9.7) that a class $\mathcal{C}$ has almost bounded flip-width if and only if it is monadically dependent. Currently, we are able to prove neither of the two implications. However, we prove the forward implication under the assumption that $\mathcal{C}$ is edge-stable (that is, excludes some half-graph as a semi-induced bipartite graph).

Theorem 9.14. *Every edge-stable, hereditary graph class of almost bounded flip-width is monadically stable.*

As far as we know, every monadically stable class might be structurally nowhere dense. This is conjectured in [NdMP+21, Conjecture 6.1]. If this were true, then Theorem 9.14 and Theorem 9.11 would imply that among edge-stable graph classes, almost bounded flip-width coincides with monadically stable class (and thus with monadically dependent classes, by Fact 9.5).

To prove Theorem 9.14, we use the following result of [GMM+23].

Fact 9.15 ([GMM+23])**.** *Let $\mathcal{C}$ be a hereditary, edge-stable class of graphs. If $\mathcal{C}$ is not monadically stable then there are $r, k \geqslant 1$, a k-coloring $\widehat{\mathcal{C}}$ of $\mathcal{C}$ and a quantifier-free formula $\varphi(x, y)$ such that $\varphi(\widehat{\mathcal{C}})$ contains the exact r-subdivision of every graph.*

Proof of Theorem 9.14. Let $\mathcal{C}$ be hereditary, edge-stable, and with almost bounded flip-width. Suppose $\mathcal{C}$ is not monadically stable. Let $r, k, \widehat{\mathcal{C}}$ and $\varphi(x, y)$ be as in Fact 9.15.

Pick a number $n \geqslant 1$. Let $K_n^{(r)}$ denote the exact r-subdivision of the clique K_n. Then $\mathrm{fw}_{r+1}(K_n^{(r)}) = \Omega_r(n)$ by Proposition 5.6. By Fact 9.15, there is some $G \in \widehat{\mathcal{C}}$ such that $\varphi(G)$ is isomorphic to $K_n^{(r)}$. In particular, $|V(G)| = |V(K_n^{(r)})| = O_r(n^2)$.

As φ is quantifier-free, by Theorem 7.2, we have $\mathrm{fw}_{r+1}(\varphi(G)) \leqslant \mathrm{fw}_{r+1}(G)$, and altogether:

$$\Omega_r(n) \leqslant \mathrm{fw}_{r+1}(K_n^{(r)}) \leqslant \mathrm{fw}_{r+1}(\varphi(G)) \leqslant \mathrm{fw}_{r+1}(G) = O_{r,\varepsilon}((n^2)^\varepsilon),$$

where the last inequality holds for every fixed $\varepsilon > 0$ since $\mathcal{C}$ (and therefore, $\widehat{\mathcal{C}}$) has almost bounded flip-width. Setting $\varepsilon = 1/4$, we get that $\Omega_r(n) = O_r(n^{1/2})$ for all n, which is absurd. Hence, $\mathcal{C}$ is monadically stable. □

9.5 Classes of ordered graphs of almost bounded flip-width

Recall that flip-width is defined for arbitrary binary relational structures (see Appendix B.4), and there is a variant of flip-width tailored to ordered graphs, defined in Section 6. It follows from Lemma 6.2, that a class $\mathcal{C}$ of ordered graphs has almost bounded flip-width if and only if for every r, we have $\mathrm{fw}_r^<(G) \leqslant |G|^{o(1)}$, for all $G \in \mathcal{C}$.

In this section, we prove that for hereditary classes of ordered graphs, almost bounded flip-width coincides with bounded twin-width, as well as with bounded flip-width (by Theorem 6.3).

Theorem 9.16. *A hereditary class $\mathcal{C}$ of ordered graphs has almost bounded flip-width if and only if $\mathcal{C}$ has bounded twin-width.*

We use the following consequence of the main result of [BGOdM+22], see [BGdM+21, Lemma 40].

Denote $[n] := \{1, \ldots, n\}$. We view $[n]^2$ as a totally ordered set, equipped with the lexicographic order $<_{\mathrm{lex}}$ on pairs.

Let $G = (V, E, <)$ be an ordered graph and $A, B \subseteq V(G)$. Fix a parameter $s \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$. Say that A and B form a *s-pattern* in G of order n^2 if $|A| = |B| = n^2$, and the following condition holds. Write $A = \{a_{ij} \mid (i,j) \in [n^2]\}$ and $B = \{b_{ij} \mid (i,j) \in [n^2]\}$, so that for $(i,j), (i',j') \in [n]^2$ with $(i,j) <_{\mathrm{lex}} (i',j')$ we have that $a_{ij} < a_{i'j'}$ and $b_{ij} < b_{i'j'}$. Then for $(i,j), (i',j') \in [n]^2$, a_{ij} and $b_{i'j'}$ are adjacent in G if and only if the following condition holds.

- $(i,j) = (i',j')$, if s is $=$,
- $(i,j) \neq (i',j')$, if s is $\neq$,
- $i = i', j \leqslant j'$, if s is $\leqslant_l$,
- $i = i', j \geqslant j'$, if s is $\geqslant_l$,
- $i \leqslant i', j = j'$, if s is $\leqslant_r$,
- $i \geqslant i', j = j'$, if s is $\geqslant_r$.

Fact 9.17 ([BGOdM+22])**.** *Let $\mathcal{C}$ be a class of ordered graphs of unbounded twin-width. Then there is $s \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ such that for every $n \geqslant 1$, there is some graph $G \in \mathcal{C}$ and sets A and B that form an s-pattern of order n^2.*

Proof. Follows from [BGdM⁺21, Lemma 40], by considering the (s, σ)-matching, where σ is the permutation of $[n]^2$ (ordered by $<_{\text{lex}}$) such that $\sigma((i,j)) = (j,i)$ for $(i,j) \in [n]^2$. □

Proof sketch of Theorem 9.16. By Theorem 6.1, if $\mathcal{C}$ has bounded twin-width then $\mathcal{C}$ has bounded flip-width, hence it has almost bounded flip-width.

Conversely, suppose $\mathcal{C}$ is a hereditary graph class with unbounded twin-width. We prove that $\mathcal{C}$ does not have almost bounded flip-width.

Let $s \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ be as in Fact 9.17.

Fix $n \geqslant 1$. By Fact 9.17 there is an ordered graph $G_n \in \mathcal{C}$ and sets $A, B \subseteq V(G_n)$ of size n^2, that form an s-pattern of order n^2 in G. As $\mathcal{C}$ is hereditary, we may assume that $V(G_n) = A \cup B$, so in particular $|V(G_n)| \leqslant 2n^2$.

Let $\mathcal{L}$ be the partition of A into n intervals of size n each, and similarly, let $\mathcal{R}$ be the partition of B into n intervals of size n each. It is not difficult to verify (by considering the three cases depending on $s \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$, and using the definition of an s-pattern of order n^2) that $\mathcal{L}$ and $\mathcal{R}$ form an n-rich division of G_n (see definition before Lemma 6.5).

By Lemma 6.5, for all $n \geqslant 1$, we have that $\mathrm{fw}_1^<(G_n) \geqslant n$. As $|G_n| \leqslant 2n^2$, we have that $\mathrm{fw}_1^<(G_n) \geqslant (|G_n|/2)^{1/2}$. Since $G_n \in \mathcal{C}$ for all $n \geqslant 1$, it cannot be that $\mathrm{fw}_1^<(G) \leqslant |G|^{o(1)}$ holds for all $G \in \mathcal{C}$. Hence, $\mathcal{C}$ does not have almost bounded flip-width. □

To summarize, in Section 9 we have defined graph classes of almost bounded flip-width. We conjecture that they coincide with monadically dependent graph classes. We provide the following evidence towards this conjecture. We have shown that, when restricted to weakly sparse graph classes, almost bounded flip-width coincides with nowhere denseness (and therefore, with monadic dependence). And when restricted to edge-stable graph classes, almost bounded flip-width generalizes structurally nowhere denseness, and is generalized by monadic stability. As it is conjectured that structurally nowhere dense classes coincide with monadically stable classes, this would imply that for edge-stable graph classes, almost bounded flip-width coincides with monadic stability (and therefore, with monadic dependence). Finally, we show that for classes of ordered graphs, almost bounded flip-width coincides with bounded flip-width, and with bounded twin-width, and therefore, with monadic dependence.

10 Discussion

In this section, we discuss some conjectured relationships between various notions defined in this paper, and other known notions. We speculate on possible routes towards proving some conjectures.

10.1 Obstructions to small flip-width

The results of this paper indicate that the flip-width parameters may be the sought dense analogue of generalized coloring numbers, and that classes of bounded flip-width may be the sought dense analogue of classes with bounded expansion. What is currently missing to complete this picture, is an analogue of the key result of Sparsity theory, which is a min-max theorem that relates, on one hand,

explicit descriptions of winning strategies of the robber, and on the other hand, explicit descriptions of winning strategies for the cops in the cop-width game (see Theorem 3.5). In the sparse case, the former are obstructions to having small r-admissibility numbers, and can be obtained from $\leqslant r$-subdivisions of graphs with large minimum degree. Finding an analogous notion in the dense case seems to be a major challenge. It seems plausible that finding such notions is related to the question of efficiently approximating the flip-width parameters. This goal can be formalized as follows.

Goal 10.1. *Obtain an fpt approximation algorithm for radius-r flip-width: an algorithm running in time $f(r,k)\cdot n^c$, for some constant c and function $f\colon \mathbb{N}\times\mathbb{N}\to\mathbb{N}$, which given an n-vertex graph G and numbers r, k, either concludes that $\mathrm{fw}_r(G)>k$, or that $\mathrm{fw}_r(G)<f(r,k)$.*

A first step in this direction is obtained by Theorem 8.7, which achieves an XP approximation algorithm, rather than an fpt algorithm.

A related, but less concrete goal is the following.

Goal 10.2. *Describe explicit* forbidden weak obstructions at depth r and density ℓ, *such that there are functions $f,g\colon \mathbb{N}\times\mathbb{N}\to\mathbb{N}$ such that for all $r,\ell\geqslant 1$, every graph G with $\mathrm{fw}_r(G)\geqslant f(r,\ell)$ contains such an obstruction of density ℓ as an induced subgraph, and conversely, no graph G which contains such an obstruction of density $g(r,\ell)$ as an induced subgraph satisfies $\mathrm{fw}_r(G)\geqslant \ell$.*

One attempt at formalizing what we mean by "explicit", is by requiring that the following holds: for every fixed r there is a formula $\varphi_r(x,y)$ (possibly involving colors) and a function $f_r\colon \mathbb{N}\to\mathbb{N}$ such that if G is is a forbidden weak obstruction at depth r and density at least $f_r(\ell)$ then there is some k-coloring of G such that $\varphi_r(G)$ contains an r-subdivision of a graph with minimum degree at least ℓ as an induced subgraph.

The following question suggests a path towards achieving Goal 10.2.

Question 10.3. *Fix $r\geqslant 1$. Is it the case that $\mathrm{fw}_r(\mathcal{C})<\infty$ if and only if there is a $k\in\mathbb{N}$ such that no graph $G\in\mathcal{C}$ contains a (r,k,k)-hideout (see Definition 4.15)?*

Hideouts are not induced subgraphs, as posited in Goal 10.2, but this could serve as a starting point.

The following notions are defined in [GPT21]. A *transduction ideal* is a property of graph classes that is preserved under first-order transductions: if a class $\mathcal{C}$ has the considered property, and $\mathcal{D}$ transduces in $\mathcal{C}$, then $\mathcal{D}$ has the property, too. The *dense analogue of bounded expansion* is the largest transduction ideal such that all weakly sparse classes that belong to it have bounded expansion.

By Corollary 7.3, classes of bounded flip-width form a transduction ideal, and by Theorem 5.2, they are contained in the dense analogue of bounded expansion. We conjecture that bounded flip-width is exactly the dense analogue of bounded expansion, according to the above definition. This is equivalently stated as follows.

Conjecture 10.4. *Let $\mathcal{C}$ be a class that does not have bounded flip-width. Then $\mathcal{C}$ transduces a weakly sparse class $\mathcal{D}$ which does not have bounded expansion.*

A road towards proving Conjecture 10.4 is through Goal 10.2. Indeed, if $\mathcal{C}$ contains obstructions at a fixed depth $r\geqslant 1$ and of arbitrarily large density ℓ, then according to the requirement given below Goal 10.2, $\mathcal{C}$ transduces a class $\mathcal{D}$ which

contains r-subdivisions of arbitrarily dense graphs, and therefore $\mathcal{D}$ does not have bounded expansion.

10.2 Model checking

Recall that Conjecture 1.1 predicts that for a hereditary class $\mathcal{C}$, first-order model-checking is fixed-parameter tractable (fpt) on $\mathcal{C}$ if and only if $\mathcal{C}$ is monadically dependent. Both implications are open. We now discuss approaching the backward implication, giving the upper bound: that model-checking is fpt on every monadically dependent graph class.

The following is by Theorem 5.2 and Theorem 6.3, and the results of [DKT13] and [BGOdM$^+$22], respectively.

Corollary 10.5. *Let $\mathcal{C}$ be either a weakly sparse graph class of bounded flip-width, or a class of ordered graphs of bounded flip-width. Then the model checking problem for first-order logic is fpt on $\mathcal{C}$.*

The question whether the model checking problem is fpt on every class $\mathcal{C}$ of graphs that has bounded twin-width, remains open. (The result of [BKTW20] says that this is the case, if the input graph G is given together with its contraction sequence. In some special cases, such a contraction sequence can be efficiently computed.) As bounded twin-width implies bounded flip-width, we do not know whether the model checking problem is fpt on every class $\mathcal{C}$ of graphs that has bounded flip-width. However, we believe that approximating the radius-r flip-width of a given graph G, might be easier than approximating its twin-width. This is indicated for example by Theorem 8.7, which gives an XP approximation algorithm for flip-width, while no such algorithm for twin-width is known. Still, an analogue of the model-checking result of [BKTW20] is missing, and is posed as the following goal.

Goal 10.6. *Devise an efficient representation of winning strategies for the cops such that for a fixed formula φ there is a number r, such that given an efficient representation of a winning strategy for the cops in the radius-r flip-width game on a given graph G, allows to efficiently check whether φ holds in G.*

In particular, Goal 10.6 combined with Goal 10.1 could allow to solve the model-checking problem on classes of bounded flip-width (and also on classes of bounded twin-width, without providing an fpt approximation algorithm for twin-width).

An extension of Goal 10.6 to monadically dependent classes, with the hope of confirming Conjecture 1.1, could lead through Conjecture 9.7, which characterizes the hereditary monadically dependent classes as those with almost bounded flip-width. This, however, remains very speculative, and most likely requiring further insights, on top of the ones needed to achieve Goal 10.6 and to solve Conjecture 9.7.

10.3 Stable classes of bounded flip-width

It is conjectured [GPT21, Conjecture 6] that the monadically stable classes in the dense analogue of bounded expansion have structurally bounded expansion, i.e., are transductions of classes with bounded expansion. As classes of bounded flip-width are contained in the dense analogue of bounded expansion, this would imply the following.

Conjecture 10.7. *Every monadically stable class with bounded flip-width has structurally bounded expansion.*

Conjecture 10.7 would generalize the main result of [GPT22, GPT21], which states that every monadically stable class of bounded twin-width has structurally bounded expansion. The conjunction of Conjecture 10.4 and Conjecture 10.7 implies the following duality statement for edge-stable graph classes.

Conjecture 10.8. *For every edge-stable graph class $\mathcal{C}$, either $\mathcal{C}$ is a transduction of a class with bounded expansion, or $\mathcal{C}$ transduces a weakly sparse class with unbounded expansion.*

Conjecture 10.4 and Conjecture 10.7 imply Conjecture 10.8: If $\mathcal{C}$ is not monadically stable, then $\mathcal{C}$ is not monadically dependent (by Fact 9.5), and hence transduces every weakly sparse class of graphs. If $\mathcal{C}$ is monadically stable and has unbounded flip-width, then $\mathcal{C}$ transduces a weakly sparse class that does not have bounded expansion by Conjecture 10.4. Finally, if $\mathcal{C}$ is monadically stable and has bounded flip-width, then $\mathcal{C}$ transduces in a class of bounded expansion, by Conjecture 10.7.

10.4 Restrictions of the flip-width game

Restricted classes of bounded flip-width can be defined by imposing additional constraints on the flip-width game. We consider the following restrictions[8] on the strategy of the cops in the flip-width game. We say that the strategy of the cops is:

blind if their move does not depend on the current position of the robber (but may depend on the graph and on the round number);

positional if their move is only based on the current position of the robber (and not on the past moves, nor the round number);

of bounded depth if there is a bound ℓ such that the cops win within ℓ rounds;

branching-blind if in each move, the cops propose a partition $\mathcal{P}$ of the vertex set, and bases his move only on the part of $\mathcal{P}$ which is occupied by the robber, and furthermore, ensure that the robber will remain in that part until the end of the game (otherwise, they loose).

Each of those variants of the game defines a variant of the flip-width parameter, and the related classes. Thus, a class $\mathcal{C}$ has bounded X flip-width, for $X \in \{$*blind, branching blind, positional*$\}$, if for every radius r, there is a k such that the cops win the flip-width game of radius r and width k on every graph $G \in \mathcal{C}$, using a strategy with property X. Similarly, $\mathcal{C}$ has *bounded flip-depth* if for every $r \in \mathbb{N}$ there are k, ℓ such that for every $G \in \mathcal{C}$, the cops have a strategy in the flip-width game of radius r and width k, that wins in at most ℓ rounds.

We may also consider the limit case of the above games, for $r = \infty$. Define classes with *blind ∞-flip-width*, *positional ∞-flip-width*, *∞-flip-depth*, *branching-blind ∞-flip-width*, analogously as above, but where we take $r = \infty$.

It turns out that these notions relate to notions and conjectures that have been studied earlier, as we discuss below. We start with the following observation.

Proposition 10.9. *All the above properties are preserved under first-order transductions, and under* CMSO *transductions for the ∞ variants.*

[8] I thank Rose McCarty and Pierre Ohlmann for suggesting the positional and bounded depth variants.

Proposition 10.9 follows by observing that in the proof of Theorem 7.2 and Theorem 7.6, the strategy on $\varphi(G)$ is obtained by transferring a winning strategy from G (see Section 7.2). And transferring a strategy with one of the listed properties, results in a strategy with the same property.

Infinite radius In the limit case of radius ∞, the picture is quite well-understood, thanks to the following.

Theorem 10.10. *Let $\mathcal{C}$ be a class of graphs. Then:*

1. *$\mathcal{C}$ has branching-blind ∞-flip-width if and only if $\mathcal{C}$ has bounded clique-width;*
2. *$\mathcal{C}$ has blind ∞-flip-width if and only if $\mathcal{C}$ has bounded linear clique-width;*
3. *$\mathcal{C}$ has bounded positional ∞-flip-width, if and only if $\mathcal{C}$ has bounded ∞-flip-depth, if and only if $\mathcal{C}$ has bounded shrubdepth.*

A class $\mathcal{C}$ has *bounded shrubdepth* [GHN+12] if and only if it transduces in a class of trees of bounded depth.

Sketch of proof. (1). For the forward implication, observe that the strategy presented in the proof of Theorem 4.13 (see Appendix D) is branching blind. For the backward implication, note that if $\mathcal{C}$ has unbounded clique-width then by Theorem 4.13 it has unbounded flip-width, so in particular, it has unbounded branching-blind flip-width.

(2). For the forward implication, observe that the strategy presented in Example 4.7, for half-graphs, is blind. Therefore, the class of half-graphs has bounded blind flip-width. It is well-known that every class $\mathcal{C}$ of bounded linear clique-width is a CMSO transduction of the class of half-graphs (or of the class of finite total orders). By Proposition 10.9, $\mathcal{C}$ has bounded blind flip-width.

Conversely, if a class $\mathcal{C}$ has unbounded linear clique-width, then it CMSO transduces the class of all trees (this follows from [HJMW20, DT17] and from [CiO07]). It is not difficult to verify that the class of trees does not have bounded blind ∞-flip-width. Again by Proposition 10.9, this implies that $\mathcal{C}$ has unbounded blind ∞-flip-width.

(3). First observe that in a tree of depth at most d, the cops have a positional winning strategy in the radius-∞ flip-width game, that wins in at most d rounds (this is essentially the same as strategy in Example 3.3). Hence, every class $\mathcal{D}$ of trees of bounded depth has bounded positional ∞-flip-width, and bounded ∞-flip-depth. By Proposition 10.9, the same holds for every class that transduces in $\mathcal{D}$, hence for all classes of bounded shrubdepth.

Conversely, if a class $\mathcal{C}$ has unbounded shrubdepth, then it CMSO-transduces the class of all paths, by [KMOW21]. It is not difficult to verify that the class of all paths does not have bounded positional ∞-flip-width, and does not have bounded ∞-flip-depth. Proposition 10.9 implies that $\mathcal{C}$ has unbounded positional ∞-flip-width, and unbounded ∞-flip-depth. □

Blind case We now move the case of classes of bounded blind flip-depth. We start with some examples.

Example 10.11. The winning strategy on half-graphs described in Example 4.7 is blind. Hence, the class of half-graphs has bounded blind flip-width. Classes of bounded blind flip-width include classes of bounded pathwidth, and more

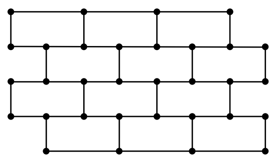

Figure 2: A wall graph

generally, classes of bounded linear rank-width (or linear clique-width), by Theorem 10.10.

Example 10.12. On the other hand, the strategies described in Example 3.3 and in Example 4.8, for trees and for comparability graphs of trees, respectively, are not blind. One can show that those classes do not have bounded blind flip-width. A more general statement is given below.

By a *subdivision* of a graph G we mean any graph obtained from G by replacing each edge by a path of arbitrary length. We omit the proof of the following proposition.

Proposition 10.13. *Let $\mathcal{C}$ be a class that contains a subdivision of every cubic tree. Then $\mathcal{C}$ does not have bounded blind flip-width.*

[GPT21, Conjecture 3] predicts that every class of unbounded linear clique-width transduces some class that contains a subdivision of every cubic tree. This would imply the following, seemingly weaker conjecture.

Conjecture 10.14. *A class $\mathcal{C}$ has bounded blind flip-width if and only if $\mathcal{C}$ has bounded linear clique-width.*

As bounded linear clique-width is equivalent to bounded blind ∞-flip-width by Theorem 10.10, Conjecture 10.14 states that bounded blind flip-width collapses to bounded blind ∞-flip-width.

Branching-blind case We now discuss the branching-blind case. By Theorem 10.10, every class of bounded clique-width has bounded branching-blind flip-width.

A wall is a graph as depicted in Figure 2. We omit the proof of the following proposition.

Proposition 10.15. *Let $\mathcal{C}$ be a class that contains a subdivision of every wall. Then $\mathcal{C}$ does not have bounded branching blind flip-width.*

[GPT21, Conjecture 4] predicts that every class of unbounded clique-width transduces some class that contains a subdivision of every wall. This would imply the following, weaker conjecture.

Conjecture 10.16. *A class $\mathcal{C}$ has bounded branching-blind flip-width if and only if $\mathcal{C}$ has bounded clique-width.*

As bounded clique-width is equivalent to bounded branching-blind ∞-flip-width by Theorem 10.10, Conjecture 10.14 states that bounded branching-blind flip-width collapses to bounded branching-blind ∞-flip-width.

Positional case Let us now look at some examples related to bounded positional flip-width.

Example 10.17. Every class with bounded expansion has bounded positional flip-width. This follows from the proof of the upper bound in Theorem 3.5, as the cops' strategy in the radius-r cop-width game is to occupy (or isolate) the vertices that are $2r$-weakly reachable from the vertex that is currently occupied by the robber. Clearly, this is a positional strategy.

By Proposition 10.9 we get the following.

Corollary 10.18. *Every class with structurally bounded expansion has bounded positional flip-width.*

On the other hand, we have:

Example 10.19. Half-graphs have unbounded positional flip-width.

It follows that every class $\mathcal{C}$ with bounded positional flip-width is edge-stable. Since $\mathcal{C}$ is also monadically dependent by Corollary 7.4, and every monadically dependent, edge-stable class is monadically stable by Fact 9.5, we get the following.

Corollary 10.20. *Every class with bounded positional flip-width is monadically stable.*

Conjecture 10.7 therefore implies the following characterization of classes with bounded positional flip-width.

Conjecture 10.21. *A class $\mathcal{C}$ has bounded positional flip-width if and only if $\mathcal{C}$ is structurally nowhere dense.*

Bounded flip-depth Finally, we look at classes of bounded flip-depth.

Example 10.22. Classes of bounded degree have bounded flip-depth, as seen in Example 3.2. Also, classes of bounded treedepth, or more generally, classes of bounded shrubdepth, have bounded flip-depth.

A class of examples combining the above, is provided by the following notion. Fix $d, h \geqslant 1$. A *hybrid tree* of depth h and degree d is a graph G that can be obtained from a rooted tree T of depth h by:

- first adding some edges connecting siblings in T, in such a way that every vertex is adjacent in G to at most d of its siblings in T,
- afterwards, taking a subgraph of the resulting graph.

Proposition 10.23. *Fix $d, h \geqslant 1$. The class $\mathcal{H}_{d,h}$ of hybrid trees of depth h and degree d has bounded flip-depth.*

The following is a consequence of Simon's factorisation theorem [Sim90]. We omit the details.

Theorem 10.24. *Every class $\mathcal{C}$ of bounded pathwidth, and more generally, every stable class of bounded linear clique-width, transduces in $\mathcal{H}_{2,h}$, for some fixed h.*

Corollary 10.25. *Every class of bounded pathwidth, and more generally, every stable class of bounded linear clique-width has bounded flip-depth.*

We now pose two conjectures characterizing classes of bounded flip-depth in complementary ways.

Conjecture 10.26. *A class $\mathcal{C}$ has bounded flip-depth if and only if $\mathcal{C}$ transduces in $\mathcal{H}_{d,h}$, for some $d, h \geqslant 1$.*

Conjecture 10.27. *A class $\mathcal{C}$ has unbounded flip-depth if and only if $\mathcal{C}$ transduces the class $\mathcal{T}_h$ of all trees of depth h, for every fixed $h \geqslant 1$.*

The conjunction of the two conjectures implies a duality statement, that a class $\mathcal{C}$ transduces in $\mathcal{H}_{d,h}$ for some $d, h \geqslant 1$ if and only if $\mathcal{C}$ does not transduce the class $\mathcal{T}_k$ of all trees of depth k, for some $k \geqslant 1$. The following conjecture makes the following, more precise min-max statement, predicting that we can take $k = h + 1$.

Conjecture 10.28. *Fix $h \geqslant 1$. A class $\mathcal{C}$ transduces in the class $\mathcal{H}_{d,h}$ for some $d \geqslant 1$ if and only if $\mathcal{C}$ does not transduce the class $\mathcal{T}_{h+1}$ of all trees of depth $h + 1$.*

The case $h = 1$ of this conjecture holds by the following duality result of Braunfeld, Nešetřil, Ossona de Mendez, and Siebertz [BNdMS22a, Theorem 7.3].

Fact 10.29. *A class $\mathcal{C}$ of graphs transduces in a class of graphs of bounded degree if and only if $\mathcal{C}$ does not transduce the class of all star forests.*

The class of all star forests is transduction-equivalent (each class transduces the other) with the class of all trees of depth 2, and the class $\mathcal{H}_{d,1}$ is transduction-equivalent with the class of graphs of degree at most d, so Conjecture 10.28 in the case $h = 1$ is a rephrasing of Fact 10.29.

We finish with the following conjecture, relating classes of bounded flip-width with the complexity of the model-checking problem for first-order logic.

Say that a graph class $\mathcal{C}$ has *non-uniform elementary-fpt* first-order model-checking if there are numbers $h, c \geqslant 1$ such that for every first-order formula φ of length k there is an algorithm which determines if a given n-vertex graph $G \in \mathcal{C}$ satisfies φ in time at most

$$\underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2^k}}}}}}_{\text{height } h} \cdot n^c.$$

In the *uniform* variant, there is a single algorithm, which inputs φ and $G \in \mathcal{C}$, and has the above running time.

It is known [DGKS07] that every class of bounded degree has uniform elementary-fpt model-checking. Furthermore, the class of trees does not have uniform elementary-fpt model-checking, unless AW[$*$]=FPT, by a result of Frick and Grohe [FG04] (see also [DGKS07]), and the class of colored linear orders does not have uniform elementary-fpt model-checking unless P=NP, also by Frick and Grohe.

Together with Michał Pilipczuk, we pose the following.

Conjecture 10.30. *A class $\mathcal{C}$ of graphs has non-uniform elementary-fpt first-order model-checking if and only if $\mathcal{C}$ does not transduce the class $\mathcal{T}_h$ of trees of height h, for some $h \geqslant 1$.*

Together with Conjecture 10.27, this would characterize classes with non-uniform elementary-fpt first-order model-checking precisely as those with bounded flip-depth.

Appendix

A Variants of the cop-width parameter

A.1 A variant without announced moves

Fix parameters $k, r \in \mathbb{N}$, and consider the variant of the Cops and Robber game in which there are k cops, and a robber with speed r, and in each round, first the cops pick a set A of k chosen vertices of the graph, and then the robber may either stay in his last position v, if $v \notin A$, or moves to any vertex u via a path $v = v_0, \ldots, v_i = u$ of length at most r such that $v_1, \ldots, v_i \notin A$. If he cannot do so, the cops win the game, and if the robber can evade the cops forever, then he wins the game. Denote the smallest number k for which the cops have a winning strategy on a graph G by $\mathrm{copwidth}'_r(G)$.

Those parameters essentially appear in the work [RT08, LPPT20]. The paper [RT08] considers a variant of the game in which the robber is *lazy*, that is, does not move unless a cop is placed at his location, whereas the [LPPT20] considers a variant where the cops occupy edges instead of vertices, and the robber never remains put. Analogues of the next notion and lemma also appear in those papers.

Call a set U of vertices of a graph G a (k, r)-*hideout* if for every $v \in U$ and set $A \subseteq V(G) - \{v\}$ with $|A| < k$, there is some path from v to $U - \{v\}$ of length at most r in $G - A$.

Lemma A.1. *Fix numbers $k, r \in \mathbb{N}$ and a graph G. The following conditions are equivalent:*

1. $\mathrm{copwidth}'_r(G) \leqslant k$,
2. *G has no (k, r)-hideout,*
3. *there is a total order on $V(G)$ such that for every $v \in V$ there is some set $A \subseteq V(G) - \{v\}$ with $|A| < k$ such that there is no path from v to any vertex $w < v$ in $G - A$.*

Proof. We show that if U is a (k, r)-hideout in G, then there is a winning strategy for the robber in the game of radius r and width k corresponding to the $\mathrm{copwidth}'_r$ parameter. As a first move, the robber picks an arbitrary vertex $v \in U$, and we show that he may always remain in the set U. In each round, when the cops place the cops on a set A of at most k vertices, then robber stays put if his current position v is not in A, and otherwise, as $|A - \{v\}| < k$, the robber is moved to any vertex $u \in U - A$ that is connected by a path of length at most r from v in $G - A$.

Suppose that G has no (k, r)-hideout. Start with $U = V(G)$ and $\bar{w}$ being the empty sequence. As long as U is nonempty, pick any vertex $u \in U$ such that there is some set $A \subseteq V(G) - \{v\}$ with $|A| < k$ such that there is no path from v to any vertex $w \in U$ in $G - A$. Such a vertex exists, since U is not a (k, s)-hideout. Remove u from U and prepend it to $\bar{w}$, and repeat. Once U becomes empty, the sequence $\bar{w}$ gives a total order on $V(G)$ satisfying the required condition.

We show how to turn a total order as in condition (3) into a winning strategy for the cops in the cop-width game of radius r and width w. When the robber is occupying a vertex v, the cops pick any set $A \subseteq V(G) - \{v\}$ with $|A| < k$ such that there is no path from v to any vertex $w < v$, and places the cops on all the vertices of $A \cup \{v\}$. Then the robber needs to move right in the order, so eventually he will lose. $\square$

We now relate the parameter $\mathrm{copwidth}'_r$ with generalized coloring numbers. The *r-strong coloring number* of a graph G, denoted $\mathrm{scol}_r(G)$ is the smallest number k with the following property. There is a total order $<$ on $V(G)$ such that every for vertex v, there are at most k vertices $w < v$ that can be reached from v by a path of length at most r. Clearly, $\mathrm{adm}_r(G) \leqslant \mathrm{scol}_r(G) \leqslant \mathrm{wcol}_r(G)$.

Lemma A.2. *For every $r \in \mathbb{N}$ and graph G the following inequalities hold:*

$$\mathrm{adm}_r(G) + 1 \leqslant \mathrm{copwidth}'_r(G) \leqslant \mathrm{scol}_r(G) + 1.$$

Proof. We prove the first inequality, by showing that $\mathrm{adm}_r(G) \geqslant k-1$ implies $\mathrm{copwidth}'_r(G) > k$. Indeed, by $\mathrm{adm}_r(G) \geqslant k-1$ there is a set $U \subseteq V(G)$ such that for every $v \in U$ there are k paths of length at most r from v to $U - v$, which are vertex-disjoint apart from v. In particular, no set $A \subseteq V(G) - \{v\}$ with $|A| < k$ hits all of those k paths. Hence, U is a (k,r)-hideout, and $\mathrm{copwidth}'_r(G) > k$ by Lemma A.1.

For the second inequality, suppose $\mathrm{scol}_r(G) \leqslant k$ and let $<$ be a total order witnessing it. The strategy of the cops in the cop-width' game with radius r is as follows: whenever the robber occupies a vertex v, then move the cops to the vertex v and all the vertices $w < v$ that are reachable from v by a path of length at most r. Then the robber needs to move some vertex $w > v$, and eventually the cops win. $\square$

Also note that $\mathrm{copwidth}'_r(G) \leqslant \mathrm{copwidth}_r(G)$, since a winning strategy in the cop-width game is also a winning strategy in the cop-width' game.

Corollary A.3. *The following conditions are equivalent for a graph class $\mathcal{C}$:*

1. *$\mathcal{C}$ has bounded expansion,*
2. $\mathrm{copwidth}_r(\mathcal{C}) < \infty$, *for every $r \in \mathbb{N}$,*
3. $\mathrm{copwidth}'_r(\mathcal{C}) < \infty$, *for every $r \in \mathbb{N}$.*

A.2 Isolation game

Consider the following variant of the flip-width game. The *isolation-width* game with radius $r \in \mathbb{N} \cup \{\infty\}$ and width $k \in \mathbb{N}$, $k \geqslant 1$, is played on a graph G. In round i of the game we have a set $S_i \subseteq V(G)$ with $|S_i| \leqslant k$, which are the new positions of the cops declared by the cops, and the current position $v_i \in V(G)$ of the robber. Initially, $S_0 = \emptyset$ and v_0 is a vertex of G chosen by the robber. In round $i > 0$, the cops announce a set $S_i \subseteq V(G)$ of next positions of the cops with $|S_i| \leqslant k$, that will be put into effect momentarily. The robber, knowing S_i, moves to a new vertex v_i by following a path of length at most r from v_{i-1} to v_i that avoids the *previous* cop positions S_{i-1}. The game terminates when $v_i \in S_i$. Write $\mathrm{iw}_r(G)$ for the smallest number k such that the cops have a winning strategy in the isolation game with radius r and width k.

Lemma A.4.

$$\mathrm{iw}_r(G) \leqslant \mathrm{copwidth}_r(G) \leqslant 2\,\mathrm{iw}_r(G).$$

Proof sketch. A winning strategy in the cop-width game of radius r and width k can be translated into a winning strategy in the isolation game of radius r and

width k: when in the cop-width game the cops are directed to their new set of positions X, in the isolation game the cops define X as their next positions. In particular, $|X| \leqslant k$, and a response of the robber in the isolation game is a valid response in the cop-width game.

Conversely, a winning strategy in the isolation game of radius r and width k can be translated into a winning strategy in the cop-width game of radius r and width $2k$: when in the isolation game the cops declare the new set X of cops positions and Y is the previous set of positions, in the cop-width game the cops define $X \cup Y$ as the next positions of the cops. In particular, $|X \cup Y| \leqslant 2k$, and a response of the robber in the cop-width game is a valid response in the isolation game. $\square$

B Flip-width

B.1 Bipartite variants

Let G be a bipartite graph. We define the parameter $\mathrm{bfw}_r(G)$ analogously to $\mathrm{fw}_r(G)$, but the flips played by the cops are now *bipartite* flips: flips between two subsets of opposite parts of G. For the parameter $\mathrm{bfw}_r^{\mathrm{part}}(G)$, we only consider partitions of $V(G)$ that refine the bipartition of $V(G)$, and measure its size by the maximum, over the two parts of G, of the number of parts of $\mathcal{P}$ that are contained in it. By definition, we have the following.

Lemma B.1. *Let G be a bipartite graph and $r \in \mathbb{N} \cup \infty$. Then $\mathrm{fw}_r(G) \leqslant 2\,\mathrm{bfw}_r(G)$.*

Let G be a graph and $X, Y \subseteq V(G)$, and let $G[X,Y]$ be the bipartite graph with parts X, Y and edges xy such that $x \in X, y \in Y, xy \in E(G)$.

Lemma B.2. *Let G be a graph and $X, Y \subseteq V(G)$, Then $\mathrm{bfw}_r^{\mathrm{part}}(G[X,Y]) \leqslant \mathrm{fw}_r^{\mathrm{part}}(G)$.*

Proof. We convert a winning strategy for the cops on G to a winning strategy on $H := G[X,Y]$. If the cops play a k-partite flip G' of G, and $\mathcal{P}$ is the corresponding partition of $V(G)$, then consider the bi-partition $\mathcal{Q}$ of $X \uplus Y$ $\mathcal{Q} := \{P \cap X, P \cap Y \mid P \in \mathcal{P}\}$, and the $\mathcal{Q}$-flip H' of H, in which two parts of $R, S \in \mathcal{Q}$ are flipped if and only if the two unique parts of $A, B \in \mathcal{P}$ such that $A \cap X = R$ and $B \cap Y = S$, are flipped in the $\mathcal{P}$-flip producing G' from G. The key property of this construction is that any path in H', starting at some vertex u and ending at a vertex v, determines a path in G' of the same length, starting at (a copy of) u and ending at (a copy of) v. In particular, if the robber moves to a vertex v' along a path of length at most r in the previous flip of H, then this induces a path of length at most r in the previous flip of G. Hence, if the cops win in G, then they also win in H. $\square$

B.2 Excluding a $K_{t,t}$

Lemma B.3. *Fix a graph G that excludes $K_{t,t}$ as a subgraph, where $t > 1$. Let $S \subseteq V(G)$ and let $\mathcal{P}_S$ be the partition of $V(G)$ that partitions S into singletons and vertices in $v \in V(G) - S$ according to $N(v) \cap S$. Then $|\mathcal{P}_S| \leqslant |S|^t$.*

Proof. Fix a set $S \subseteq V(G)$ and let $s = |S|$.

To every vertex $v \in V(G)$ with $|N(v) \cap S| \geqslant t$ assign any set $A \subseteq N(v) \cap S$ with $|A| = t$. Then every set $A \subseteq S$ with $|A| = t$ is assigned to at most $t-1$

vertices v, since otherwise we have a $K_{t,t}$ as a subgraph of G. It follows that $\{N(v) \cap S \mid v \in V(G), |N(v) \cap S| \geqslant t\}$ has at most $(t-1) \cdot \binom{s}{t}$ elements. On the other hand, $\{N(v) \cap S \mid v \in V(G), |N(v) \cap S| < t\}$ has at most s^{t-1} elements. Altogether, $\{N(v) \cap S \mid v \in V(G)\} \uplus S$, which is in bijection with the partition $\mathcal{P}_S$ of $V(G)$, has at most s^t elements. □

B.3 2VC-dimension

Corollary (5.7)**.** *If G is the exact 1-subdivision of an n-clique, then $\mathrm{fw}_2(G) > (n-1)/4$. Furthermore, for every graph G, $\mathrm{2VCdim}(G) \leqslant 8\,\mathrm{fw}_2(G) + 2$.*

Proof. The first part follows immediately from Proposition 5.6, for $r = 2$ and $k = (n-1)/4$. We prove the second part. Let $V = V(G)$, and consider the bipartite graph $G[V, V]$. Suppose $\mathrm{2VCdim}(G) \geqslant k$. Then $G[V, V]$ contains the 1-subdivision of K_k as an induced subgraph. By the first part of Corollary 5.7, we have that $\mathrm{fw}_2(G[V, V]) > (k-1)/4$. By Lemma B.1, we have $\mathrm{bfw}_r(G[V, V]) \geqslant \mathrm{fw}_r(G[V, V])/2 > (k-1)/8$, and by Lemma B.2, we have $\mathrm{fw}_r(G) \geqslant \mathrm{bfw}_r(G[V, V]) > (k-1)/8$. □

B.4 Flip-width of binary structures

We extend the definition of flip-width to structures equipped with one or more binary relations. To this end, we extend the notion of flips as follows. Let $R \subseteq V \times V$ be a binary relation on a set V, and let (A, B) be a pair of subsets of V. The relation $R' \subseteq V \times V$ obtained from R by *flipping* the pair (A, B) (now the order of the pair matters) is defined as $R' := R \triangle (A \times B)$, where $\triangle$ is the symmetric difference. We will apply such flips in the context of binary relational structures, as defined below.

Fix a binary relational signature Σ, that is, a signature consisting of unary and binary relation symbols only (see Section 2.4). Let B be a Σ-structure and $\mathcal{P}$ be a partition of $V(G)$. A *$\mathcal{P}$-flip* is an operation, which is specified by a Σ-structure F with vertex set $\mathcal{P}$. Applying this operation to B results in the Σ-structure B' with $V(B') = V(B)$ and relations

$$R_{B'} \quad := \quad R_B \quad \triangle \bigcup_{(P,Q) \in R_F} P \times Q,$$

for each binary relation symbol $R \in \Sigma$. The unary relation symbols are interpreted in B' in the same way as in B. By slight abuse of language, we sometimes call the structure B' a $\mathcal{P}$-flip of B.

For $k \geqslant 1$, a *k-flip* of B is a Σ-structure B' which is a $\mathcal{P}$-flip of B, for some partition $\mathcal{P}$ of $V(B)$ with $|\mathcal{P}| \leqslant k$.

Fix $r \in \mathbb{N} \cup \{\infty\}$. We now define the *flip-width game* of radius r and width k on a Σ-structure B similarly as in the case of graphs, with the following differences: in each round, the cops announce a k-flip B' of B, whereas the robber moves along a path of length at most r in the Gaifman graph of the k-flip of B that was announced in the previous round (and B in the first round). The radius-r flip-width of B, denoted $\mathrm{fw}_r(B)$ is the smallest k such that the cops win the flip-width game of radius r and width k on B.

Graphs are viewed as structures over the signature Σ consisting of a single binary relation symbol E, interpreted in a given graph G as the (symmetric, irreflexive) adjacency relation. Note that in principle, applying a $\mathcal{P}$-flip F to a graph G can result in a Σ-structure G' in which the binary relation is no longer symmetric. This is the case when the binary relation of F is not symmetric. However, for the cops, it never pays off to apply such flips, since the robber moves in the Gaifman graph of the resulting structure, so the directions of the edges are of no relevance to him.

Example B.4. Let $B = (V, <)$ be a totally ordered set, viewed as a structure over the signature $\Sigma = \{<\}$. Let $\mathcal{P}$ be a partition of V into k sets that are intervals with respect to $<$, and let $I_1, \ldots, I_k$ denote those intervals in increasing order. Consider the $\mathcal{P}$-flip of B specified by the Σ-structure F with vertices $\mathcal{P}$, with $<_F = \{(I_i, I_j) \mid 1 \leqslant i < j \leqslant k\}$. Applying the flip F results in the Σ-structure $B' = (V, <')$, where $a <' b$ if and only if $a < b$ and a and b belong to the same part of $\mathcal{P}$.

Now consider the $\mathcal{P}$-flip of B specified by the Σ-structure F' with vertices $\mathcal{P}$, with $<_F = \{(I_i, I_j) \mid 1 \leqslant i \leqslant j \leqslant k\}$ (so we also flip pairs (I_i, I_i)). Applying the flip F' results in the Σ-structure $B' = (V, <')$, where $a <' b$ if and only if $a \geqslant b$ and a and b belong to the same part of $\mathcal{P}$ (so the order is reversed and becomes reflexive).

Example B.5. The radius-∞ flip-width of a totally ordered set $B = (\{1, \ldots, n\}, <)$ is at most three. The strategy is similar as in Example 4.7, but now in the ith round, the cops apply the 3-flip F of B as in the previous example, for the partition $\mathcal{P}$ of V into the intervals $\{1, \ldots, i\}, \{i\}, \{i+1, \ldots, n\}$, removing all relations between distinct intervals.

B.5 Definable flip-width

We prove Lemma 8.2.

Lemma (8.2). *There is an algorithm that, given a graph G and numbers $k \in \mathbb{N}$ and $r \in \mathbb{N} \cup \{\infty\}$, determines whether* $\mathrm{dfw}_r(G) \leqslant k$ *in time* $n^{O(k)} \cdot 2^{O(2^k)}$.

Proof. Fix k, r, and a graph G. A *configuration* in the definable flip-width game with radius r consists of:

- a set $S \subseteq V(G)$ of size at most k, specifying the partition $\mathcal{P}$ played by the cops,
- a $\mathcal{P}$-flip of G,
- the current position of the robber.

The set of all configurations has size $O(n^{k+1} \cdot 2^{4^k})$, and the winner of the game can be computed using a fixpoint computation running in time polynomial in $n^{k+1} \cdot 2^{4^k}$. ∎

Lemma B.6 (8.3). *Fix $r \in \mathbb{N} \cup \{\infty\}$. For every graph G we have:*

$$\mathrm{fw}_r(G) = O(\mathrm{dfw}_r(G)^{\mathrm{VCdim}(G)}). \tag{14}$$

Proof. Let G be a graph and let $d = \mathrm{VCdim}(G)$ and let $k = \mathrm{dfw}_r(G)$. Consider a set $S \subseteq V(G)$ with $|S| \leqslant k$. Then the set system $(S, \{N(v) \cap S \mid v \in V(G)\})$

has VC-dimension at most d. By the Sauer-Shelah-Perles lemma, we have that $\{N(v) \cap S \mid v \in V(G)\}$ has $O(|S|^d) = O(k^d)$ elements.

It follows that every k-definable flip of G is a $O(k^d)$-flip of G. Since $\mathrm{dfw}_r(G) \leqslant k$ it follows that $\mathrm{fw}_r(G) = O(k^d)$. $\square$

C Modular partition and substitution closure

A set of vertices $X \subseteq V(G)$ in a graph X is a *module* if all vertices in X have the same neighbors outside of X. Note that in a modular partition, all the parts of the partition are modules.

Lemma C.1 (4.10). *Let G be a graph and $\mathcal{P}$ be its modular partition. Then*

$$\mathrm{fw}_r(G) \leqslant \max\left(\mathrm{fw}_r(G/\mathcal{P}), \max_{A \in \mathcal{P}} \mathrm{fw}_r(G[A]) + 2\right).$$

Proof sketch for Lemma 4.10. The strategy of the cops is as follows. First, use a strategy of width $\ell = \mathrm{fw}_r(G/\mathcal{P})$ on $G/\mathcal{P}$. Each ℓ-flip G' of $G/\mathcal{P}$ induces an ℓ-flip $\widehat{G}$ of G. Playing according to this strategy, the robber eventually reaches a vertex $A \in V(G/\mathcal{P}) = \mathcal{P}$ that is isolated in the current flip G' of $G/\mathcal{P}$. The corresponding play in G leads to a ℓ-flip $\widehat{G}$ of G such that the current vertex v of the robber is in A, and there are no edges joining A and $V(G) - A$ in $\widehat{G}$.

Since A is a module in G, the graph G_A obtained from G by removing all edges with one endpoint in A and one endpoint in $V(G) - A$, can be obtained from G by flipping A and $N(A) - A$. The cops now announce the graph G_A, and the robber is still confined to A.

Now, we use a winning the strategy of the cops in the graph $G[A]$, of width $k = \mathrm{fw}_r(G[A])$. Whenever in the game on $G[A]$ the cops announce a k-flip G' of $G[A]$, in the game on G the cops announce the graph $\widehat{G}$ such that $\widehat{G}[A] = G'[A]$, $\widehat{G}[V(G) - A] = G[V(G) - A]$, and there are no edges joining A and $V(G) - A$ in $\widehat{G}$. The graph $\widehat{G}$ is a $(k+2)$-flip of G, where the partition partitions A into k parts, according to the partition of $G[A]$ used in the k-flip G' of $G[A]$, and partitions $V(G) - A$ into two parts, $N(A) - A$ and $V(G) - (N(A) \cup A)$. It follows that playing according to this strategy, the cops will win, once the cops win in the game on $G[A]$. $\square$

Lemma (4.11). *Fix $r \in \mathbb{N} \cup \{\infty\}$ For every $r \in \mathbb{N} \cup \{\infty\}$ and graph class $\mathcal{C}$, we have*

$$\mathrm{fw}_r(\mathcal{C}^*) \leqslant \mathrm{fw}_r(\mathcal{C}) + 2.$$

In particular, if $\mathcal{C}$ has bounded flip-width, then $\mathcal{C}^*$ has bounded flip-width.

The idea is to use the winning strategy on the class $\mathcal{C}$, and confine the robber to $L(w)$, for deeper and deeper nodes of w, similarly as in the strategy given Example 3.3 for the cop-width game on trees.

Proof sketch for Lemma 4.11. Suppose $\mathrm{fw}_r(\mathcal{C}) \leqslant k$. Let $G \in \mathcal{C}^*$. Note that G has a modular partition $\mathcal{P}$ such that $G/\mathcal{P} \in \mathcal{C}^*$ and each part A of $\mathcal{P}$ induces a graph $G[A] \in \mathcal{C}^*$, or is a singleton. By repeatedly applying the argument as in the proof of Lemma 4.10, we conclude that $\mathrm{fw}_r(G) \leqslant k + 2$. $\square$

D Flip-width with infinite radius

Theorem D.1 (4.13)**.** *A class of graphs $\mathcal{C}$ has bounded clique-width if and only if* $\mathrm{fw}_\infty(\mathcal{C}) < \infty$.

Proof. We prove that every class $\mathcal{C}$ of bounded clique-width has bounded flip-width at radius-∞. The other implication was shown in Section 4.2.

Fix a number k and let $\mathcal{C}$ be a class of graphs of clique-width at most k. Then for every $G \in \mathcal{C}$ there is a rooted binary tree T with leaves $V(G)$, such that for every node w of T the set L_w of leaves that are descendants of w induces at most $O_k(1)$ distinct neighborhoods over its complement $V(G) - L_w$.

We present a strategy for the cops, for the radius $r = \infty$. The cops store a node w_i of T, maintaining the invariant that in round i, w_i is at distance exactly i from the root, and $A_i \subseteq L_{w_i}$, where A_i denotes the set of possible next moves of the robber.

Initially, w_0 is the root, and $A_0 = V(G)$, so the invariant is satisfied.

We now describe the cops strategy that maintains the invariant. Suppose we are in round i. If w_i is a leaf then the cops win, since $A_i \subseteq \{w_i\}$ by the invariant. Otherwise, partition $V(G)$ into three parts:

- B_0 – the leaves of T below the left child of w_{i-1},
- B_1 – the leaves of T below the right child of w_{i-1},
- B_2 – the remaining leaves of T.

For $j = 0, 1, 2$, let $\mathcal{Q}_j$ be the partition of the set B_j into equivalence classes of the relation of having equal neighborhoods in $V(G) - B_j$. Then $|\mathcal{Q}_j| \leqslant O_k(1)$, for $j = 0, 1, 2$. Define the partition $\mathcal{P}_i := \mathcal{Q}_0 \cup \mathcal{Q}_1 \cup \mathcal{Q}_2$. Then there is a $\mathcal{P}_i$-flip G_i of G such that for any edge $uv \in E(G_i)$, the vertices u and v belong to the same part of the partition $\{B_0, B_1, B_2\}$. The cops play the partition $\mathcal{P}_i$ and the flip G_i.

Suppose the robber responds by moving to a vertex $c_i \in A_{i-1}$. Since A_{i-1} consists of descendants of w_{i-1} by the invariant, it follows that either $c_i \in B_0$ or $c_i \in B_1$. Therefore, the set A_i, which is defined as the connected component of c_i in G_i, is either contained in B_0, or in B_1. In the first case, define w_i as the left child of w_{i-1}, and in the latter case, as the right child of w_{i-1}. Thus, the invariant is maintained.

Since T is a finite tree, after at most $|T|$ rounds w_i is a leaf, and then cops win. □

E VC-dimension

E.1 Radius one

By a similar argument as in Lemma 4.17, we prove:

Lemma (4.18)**.** *Let $b, k \in \mathbb{N}$ and let G be a bipartite graph with* $\mathrm{bfw}_1^{\mathrm{part}}(G) \leqslant k$*. Then G contains more than b mutual $2bk$-near-twins contained in a single part of G.*

To prove Theorem 4.21, we show that every bipartite graph of sufficiently large VC-dimension contains an induced subgraph with no k-near twins. First we construct such bipartite graphs. For a number $m \in \mathbb{N}$ let $\mathbb{Z}_2^m$ denote the m-dimensional vector space over the two-element field $\mathbb{Z}_2$, and for $v, w \in \mathbb{Z}_2^m$ denote $v \cdot w = \sum_{i=1}^m v_i \cdot w_i \mod 2$.

Lemma E.1. *Fix a number $m \in \mathbb{N}$ and consider the bipartite graph H whose parts are two copies of $\mathbb{Z}_2^m$, with edges vw such that $v \cdot w \neq 0$. Then H has no $(2^{m-1}-1)$-near-twins in either part.*

Proof. Denote the two copies of $\mathbb{Z}_2^m$ with V and V^*, and write v^* for the copy of a vector $v \in V$ in V^*. We claim that for all $u, v \in V$ we have $|N_H(u) \triangle N_H(v)| = 2^{m-1}$. Indeed,

$$N_H(u) \triangle N_H(v) = \{w^* \mid w \in V, w \cdot u \neq w \cdot v\} = \{w^* \mid w \in V, w \cdot (u-v) \neq 0\}.$$

Since $\{w \in V \mid w \cdot (u-v) \neq 0\}$ has 2^{m-1} elements, so does $N_H(u) \triangle N_H(v)$.

By a dual argument, $|N_H(u^*) \triangle N_H(v^*)| = 2^{m-1}$ for all $u^*, v^* \in Y$. Therefore, H_m has no pair of $(2^{m-1}-1)$-near twins in either part. □

Lemma (4.22)**.** *Let G be a graph with $\mathrm{VCdim}(G) \geqslant 2^m$, for some m. Then there are two sets X, Y such that the bipartite graph $G[X, Y]$ contains a pair of $(2^{m-1}-1)$-near-twins in either of the parts X, Y.*

Proof. Let X be a subset of $V(G)$ of size 2^m that is shattered by $\{N(v) \mid v \in V(G)\}$.

Arbitrarily identify the elements of X with the elements of $\mathbb{Z}_2^m$. Denote $v^\perp := \{u \in X \mid u \cdot v = 0\} \subseteq X$. Since X is shattered in G, for every $v \in X$ there is a vertex $v^* \in V(G)$ such that $N(v^*) \cap X = v^\perp$. Denote $Y := \{v^* \mid v \in X\}$. The function $v \mapsto v^*$ is then a bijection between X and Y, and for $v, w \in X$, v is adjacent to w^* in G if and only if $v \cdot w \neq 0$. Therefore, $G[X, Y]$ is isomorphic to the graph from Lemma E.1, and the conclusion follows from Lemma E.1. □

Proof of Theorem 4.21. We show that if $\mathrm{VCdim}(G) \geqslant 2^m$, then $\mathrm{fw}_1(G) > 2^{m-2}$. Assume $\mathrm{VCdim}(G) \geqslant 2^m$, and let $G[X, Y]$ be as in Lemma 4.22. By Lemma 4.18 (setting $b = 1$), this implies that $\mathrm{bfw}_1^{\mathrm{part}}(H) > 2^{m-2}$. By Lemma B.2, $\mathrm{fw}_1^{\mathrm{part}}(G) > 2^{m-2}$. □

F Exact subdivisions

We prove Proposition 5.6, which is repeated below.

Proposition (5.6)**.** *Fix $r \geqslant 2, k \geqslant 1$. Let G be the exact $(r-1)$-subdivision of some graph H with minimum degree at least $2rk$. Then $\mathrm{fw}_r(G) > k$.*

We first prove two lemmas.

Lemma F.1. *Fix $k, \ell, m \geqslant 1$. Let H be a bipartite graph with bipartition (L, R), in which every vertex in L has degree at least ℓ, and any two distinct vertices in L have at most m common neighbors. Let H' be a k-flip of H. Then there is a set $X \subseteq L$ with $|X| \geqslant |L| - k$, and a bijection $\pi\colon X \to X$, such that every $v \in X$ is adjacent in H' to at least $\lceil \frac{\ell-m}{2} \rceil$ vertices in $N_H(\pi(v))$.*

Proof. Let $\mathcal{P}$ be the partition of $V(G)$ with $|\mathcal{P}| \leqslant k$ such that H' is a $\mathcal{P}$-flip of H.

First consider the case when all the vertices of L are in one part A of $\mathcal{P}$. Let W be the union of the parts B of $\mathcal{P}$ that are not flipped with A in the flip that produces H' from H. Let X_1 consist of those vertices $v \in L$ such that $|N_H(v) \cap W| \geqslant \frac{\ell-m}{2}$, and let X_2 consist of the remaining vertices in L.

Every vertex in X_1 is adjacent in H' to at least $\frac{\ell-m}{2}$ vertices in $N_H(v)$ (namely, to $N_H(v) \cap W$), so we can set $\pi(v) = v$ for all $v \in X_1$.

If v and v' are two distinct vertices in X_2, then $N_{H'}(v) \supseteq N_H(v') \cap W$ and $|N_H(v') \cap W| \geqslant \frac{\ell-m}{2}$. If $|X_2| \leqslant 1$ then set $X := X_1 = L - X_2$, and let $\pi\colon X \to X$ be the identity on X. If $|X_2| > 2$, set $X := L = X_1 \cup X_2$, and let $\pi\colon X \to X$ be a permutation that maps every vertex in X_1 to itself, and acts as a cyclic permutation on the vertices in X_2. In any case, $|X| \geqslant |L| - 1$, and every $v \in X$ is adjacent in H' to at least $\lceil \frac{\ell-m}{2} \rceil$ neighbors of $\pi(v)$.

In the general case, partition L as $L = L_1 \uplus \ldots \uplus L_s$, for some $s \leqslant k$, following the partition $\mathcal{P}$ restricted to L. For each $1 \leqslant i \leqslant s$, let $H_i = H[L_i, R]$, and $H'_i = H'[L_i, R]$; then H'_i is a k-flip of H_i, and they fall into the special case considered above. Hence, for each $1 \leqslant i \leqslant s$ there are set $X_i \subseteq L_i$ with $|L_i| \geqslant |X_i| - 1$ and a bijection $\pi_i\colon X_i \to X_i$. Set $X = X_1 \uplus \ldots \uplus X_s$, and let $\pi\colon X \to X$ be such that $\pi(x) = \pi_i(x)$ for $x \in X_i$. Then X and π satisfy the required condition. □

Applying Lemma F.1 in the case $\ell = 1$ and $m = 0$, we get the following.

Corollary F.2. *Fix $k \geqslant 1$. Let M be a matching between two sets L and R, and let M' be a k-flip of M. Then M' contains, as a subgraph, a matching between all but k vertices of L, and a set of vertices of R of the same size.*

For $r, \ell \geqslant 1$, let $G_{r,\ell}$ denote the union of ℓ paths, each of length r (and with $r+1$ vertices), and in each path, call one of the vertices of degree one a *source*, and the other one a *target*. By an easy induction on $r \geqslant 1$, Corollary F.2 gives the following.

Lemma F.3. *Fix $k, r, \ell \geqslant 1$. If G' is a k-flip of $G_{r,\ell}$, then at least $\ell - rk$ target vertices of $G_{r,\ell}$ are joined by a path of length r in G' with some source vertex.*

Proof of Proposition 5.6. Fix $r \geqslant 2$ and let G be an exact $(r-1)$-subdivision of some graph H with minimum degree at least $2rk$. We aim to prove that $\mathrm{fw}_r(G) > k$. Let $P \subseteq V(G)$ denote the set of *principal vertices* of G, that is, vertices of degree larger than two; those vertices correspond to the vertices of H. We show that P forms a (r, k, k)-hideout in G. By Lemma 4.16, this implies that $\mathrm{fw}_r(G) > k$.

Let $R = N_G(P) = \bigcup_{v \in P} N_G(v)$ denote the set of neighbors of the principal vertices in G. Note that R and P are disjoint, as $r \geqslant 2$. Then $G[P, R]$ is a bipartite graph in which every vertex in P has at least $2rk$ neighbors in R, and any two vertices in P have at most one common neighbor in R (in fact, no common neighbors if $r \geqslant 2$). For each $v \in P$, there is an induced subgraph G_v of G that is isomorphic to $G_{r-1,\ell}$, for some $\ell \geqslant 2rk$ with the source vertices equal to $N_G(v)$, and where each target vertex is a principal vertex.

Let G' be a k-flip of G. Call a principal vertex $v \in P$ *good* if there is some $w \in P$ such that if v is adjacent in G' to at least rk elements of $N_G(w)$, and call v *bad* otherwise. We show that (1) there are at most k bad vertices, and (2) every good vertex has more than k principal vertices in its $(r+1)$-neighborhood in G'. It follows that the principal vertices form a (r, k, k)-hideout in G.

(1) Apply Lemma F.1 to the bipartite graph $G[P, R]$ and to $G'[P, R]$, which is a k-flip of F. Since every vertex in P has degree at least $2rk$, and any two distinct vertices in L have at most one common neighbor, and $\lceil (2rk-1)/2 \rceil = rk$, there are at most k bad principal vertices by Lemma F.1.

(2) Let $v \in P$ be a good principal vertex, and let $w \in P$ be such that some set A of rk vertices of $N_G(w)$ are adjacent to v in G'. Consider the subgraph K of G_w, consisting of vertex-disjoint paths of length $r-1$ in G joining the vertices of A with rk principal vertices. In particular, K is isomorphic to $G_{r-1,rk}$. Let $K' = G'[V(K)]$; then K' is a k-flip of K. By Lemma F.3, at least $rk-(r-1)k = k$ principal vertices are joined by a path of length $r-1$ in K' with some vertex in A. Those principal vertices are therefore at distance at most r in G' from v, so v has more than k principal vertices in its neighborhood in G'. □

G Bounded twin-width

G.1 Bounding flip-width in terms of twin-width

We prove:

Theorem (6.1). *Fix $r \in \mathbb{N}$. For every graph G of twin-width d we have:*

$$\mathrm{fw}_r(G) = 2^d \cdot d^{O(r)}.$$

In particular, every class of bounded twin-width has bounded flip-width.

We use the following result, proved in [BFLP23].

Theorem G.1. *Let G be a graph of twin-width $d \geqslant 1$, and let $A \subseteq V(G)$. Then we have:*

$$|\{N(v) \cap A \mid v \in V(G)\}| \quad \leqslant \quad 2^{d+O(\log d)} \cdot |A|.$$

This improves a previous result [BKR+22, Prz22], which gave a doubly-exponential dependency on d.

Proof. Fix an uncontraction sequence $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$ of G of red-degree d. For $v \in V(G)$ and $1 \leqslant i \leqslant n$ let $B^i_r(v) \subseteq \mathcal{P}_i$ denote the ball of radius r in the red graph of $\mathcal{P}_i$, around the part of $\mathcal{P}_i$ containing v. In particular, $|B^i_r(v)| < d^{r+1} = O_{r,d}(1)$.

We describe a strategy of the cops that guarantees that the following invariant holds after round i, for $i = 1, 2, \ldots, n$:

$$A_i \subseteq \bigcup B^i_r(c_i), \tag{15}$$

where c_i is robber's position in round i, and $A_i \subseteq V(G)$ is the set of possible positions the robber can move to in round $i+1$ of the game, that is, A_i is the ball of radius r around c_i in the graph G_i announced by the cops in round i, with $G_1 = G$ and c_1 being the initial vertex of the robber.

Note that for $i = n$, the inclusion (15) implies that $|A_n| = 1$ as in $\mathcal{P}_n$ there are no red edges and each part is a singleton. Therefore, the cops win the game after n rounds according to this strategy. We need to show how the cops can maintain the invariant (15), by playing k-flips of G for some k bounded in terms of r and d.

Before describing the strategy, we make an observation that will be useful in the inductive reasoning. It gives a description of a ball in the red graph of $\mathcal{P}_{i-1}$, in terms of the red graph of $\mathcal{P}_i$. Below, for $\mathcal{F} \subseteq \mathcal{P}_i$, the set $B^i_r(\mathcal{F}) \subseteq \mathcal{P}_i$ denotes the set of parts of $\mathcal{P}_i$ that are at distance at most r from some part in $\mathcal{F}$ in the red graph of $\mathcal{P}_i$.

Claim G.1. *Let $1 < i \leqslant n$ and let $v \in V(G)$. Then there is a set $\mathcal{F} \subseteq \mathcal{P}_i$ with $|\mathcal{F}| \leqslant d+3$, such that $B_r^{i-1}(v) \subseteq B_r^i(\mathcal{F})$.*

Proof. The family $\mathcal{F}$ consists of:

- the part of $\mathcal{P}_i$ that contains v,
- the parts A, B of $\mathcal{P}_i$ such that $A \cup B$ is a part of $\mathcal{P}_{i-1}$,
- the parts in $\mathcal{P}_i$ that are not homogeneous towards $A \cup B$ in G.

It can be easily verified that $\mathcal{F}$ satisfies the statement of the claim. □

We now describe the cops' strategy. In the first round, we have $G_1 = G$, and the invariant (15) is trivially satisfied since $\mathcal{P}_1$ has just one part, and that part contains $A_1 = B_G^r(c_1)$, regardless of the robber's choice of c_1.

Suppose that the invariant (15) is satisfied after round $i-1$, for some $1 < i \leqslant n$, so that $A_{i-1} \subseteq \bigcup B_r^{i-1}(c_{i-1})$. We describe how the cops should play to maintain invariant (15) after round i. Apply Claim G.1 to $v = c_{i-1}$, obtaining a family $\mathcal{F} \subseteq \mathcal{P}_i$ with $|\mathcal{F}| \leqslant d+3$ such that $B_r^{i-1}(c_{i-1}) \subseteq B_r^i(\mathcal{F})$. In particular, $A_{i-1} \subseteq \bigcup B_r^i(\mathcal{F})$.

Let $R = V(G) - \bigcup B_{2r}^i(\mathcal{F})$, and let $\mathcal{R}$ be the partition of R according to the equivalence relation of having the same neighborhood in the set $\bigcup B_{2r-1}^i(\mathcal{F})$. Note that we are simultaneously considering balls around $\mathcal{F}$ with radii $r, 2r-1$, and $2r$.

Claim G.2. *We have $|\mathcal{R}| = 2^d \cdot d^{O(r)}$, and every part $P \in \mathcal{R}$ is homogeneous towards all parts in $B_{2r-1}^i(\mathcal{F})$.*

Proof. Pick a set X which contains one representative of each part in $B_{2r-1}^i(\mathcal{F})$. Then $|X| \leqslant d^{O(r)}$. Note that every part $P \in \mathcal{P}_i - B_{2r}^i(\mathcal{F})$ is homogeneous towards every part $Q \in \mathcal{B}_{2r-1}^i(\mathcal{F})$. Therefore, the neighborhood of a vertex $v \in R$ in $\bigcup B_{2r-1}^i(\mathcal{F})$ is completely determined by $N(v) \cap X$. By Theorem G.1, there are at most

$$2^{d+O(\log d)}|X| \leqslant 2^d \cdot d^{O(r)}$$

different neighborhoods of vertices $v \in R$ in X, so $|\mathcal{R}| \leqslant 2^d \cdot d^{O(r)}$. □

Let $\mathcal{P}_i' = B_{2r}^i(\mathcal{F}) \cup \mathcal{R}$. Then $\mathcal{P}_i'$ is a partition of $V(G)$. We have $|B_{2r}^i(\mathcal{F})| \leqslant |\mathcal{F}| \cdot d^{2r+1} \leqslant O(d^{2r+2})$, and it follows from Claim G.2 that $|\mathcal{P}_i'| \leqslant 2^d \cdot d^{O(r)}$.

The cops play the $\mathcal{P}_i'$-flip G_i of G obtained by flipping any pair P, Q of distinct parts of $\mathcal{P}_i'$ such that the pair P, Q is complete in G.

Now, the robber makes his move, and picks a vertex $c_i \in A_{i-1}$, and we set $A_i := B_r^{G_i}(c_i)$. We now prove that the invariant (15) holds.

Claim G.3. $B_r^{G_i}(c_i) \subseteq \bigcup B_r^i(c_i)$.

Proof. Note that $c_i \in A_{i-1} \subseteq \bigcup B_r^i(\mathcal{F})$.

Let $v \in B_r^{G_i}(c_i)$ be a vertex at distance k from c_i in G_i, for some $0 \leqslant k \leqslant r$. We show by induction on k that if A is the part of $\mathcal{P}_i$ that contains v, then $A \in B_k^i(c_i)$. For $k = r$, this immediately yields the claim.

For $k = 0$ the statement holds trivially, so suppose that $k > 0$ and the statement holds for $k-1$. Let w be a vertex that is a neighbor of v and is at distance $k-1$ from c_i in G_i, and let B be the part of $\mathcal{P}_i$ that contains w. By inductive assumption,

$B \in B^i_{k-1}(c_i)$, and since $c_i \in \bigcup B^i_r(\mathcal{F})$, it follows that $B \in B^i_{r+k}(\mathcal{F})$. As $k < r$, we have that $B \in B^i_{2r-1}(\mathcal{F})$.

Since A and B are connected by an edge in G_i, it must be the case that A and B are not homogeneously connected in G. As $B \in B^i_{2r-1}(\mathcal{F})$, it cannot be that $A \in \mathcal{R}$ (since all parts in $\mathcal{R}$ are homogeneous towards all parts in $B^i_{2r-1}(\mathcal{F})$ by Claim G.2), so $A \in B^i_{2r}(\mathcal{F})$. As A and B are not homogeneously connected in G, A is a neighbor of B in the red graph of $\mathcal{P}_i$. As $B \in B^i_{k-1}(c_i)$, it follows that $A \in B^i_k(c_i)$, finishing the inductive step. □

The invariant (15) now follows from Claim G.3. In particular, if the cops continue playing this way, at the end of round n we have that $|A_n| = 1$, so the cops win. □

G.2 Ordered flip-width

We prove Lemma 6.2, repeated below.

Lemma (6.2). *Fix $r \in \mathbb{N} \cup \{\infty\}$ and an ordered graph $G = (V, E, <)$. Then*

$$\mathrm{fw}_r(G)^{\frac{1}{2}} \quad \leqslant \quad fw^<_r(G) \quad \leqslant \quad 2\,\mathrm{fw}_{3r+2}(G) + 1.$$

We first study the effects of applying flips to a set equipped with a total order.

Lemma G.2. *Let $L = (V, <)$ be a total order and let L' be a k-flip of L (as a binary structure). Then there is a set $S \subseteq V$ with $|S| \leqslant k$ such that any two vertices of $V - S$ with no vertex of S between them are at distance at most 2 in the Gaifman graph of L'.*

Proof. Let $\mathcal{P}$ be a partition of V with $|\mathcal{P}| \leqslant k$ such that L' is a $\mathcal{P}$-flip of L. Let $S = \{\max(A) \mid A \in \mathcal{P}\}$ be the set of $<$-maximal elements of each part of $\mathcal{P}$. We claim that S satisfies the required condition.

Observe first that each part $A \in \mathcal{P}$ forms a clique in the Gaifman graph of L'. Indeed, the relation $<'$ of L' restricted to A either coincides with $<$, if (A, A) is not flipped in the $\mathcal{P}$-flip producing L' from L, or coincides with $\geqslant$, if (A, A) is flipped. In any case, $<'$ is a total relation on A.

Let $a, b \in V - S$ be vertices belonging to different parts A, B of $\mathcal{P}$, with no element of S between a and b. We show that a and b are either adjacent, or have a common neighbor in the Gaifman graph of L'. By symmetry, suppose that $a < b$. If the pair (A, B) is not flipped in the $\mathcal{P}$-flip producing L' from L, then $a <' b$ in L', so a and b are adjacent in the Gaifman graph of L'.

Suppose that the pair (A, B) is flipped. Let $m_a := \max A$ and $m_b := \max B$. Clearly, $b < m_b$ and $a < m_a$. As m_a is not between a and b, we have that $b < m_a$. Since (A, B) is flipped and $m_a \not< b$ in L, it follows that $m_a <' b$ in L'. Since a and m_a are in the same part of A, they are adjacent in the Gaifman graph of L'. Therefore, m_a is a common neighbor of a and of b in the Gaifman graph of L'. □

Corollary G.3. *Let $G = (V, E, <)$ be an ordered graph, and let $G' = (V, E', <')$ be a k-flip of G, in the sense of binary structures. Then there is a $(2k+1)$-cut-flip $G'' = (V, E', \sim')$ of G, such that for all $u, v \in V$, if u and v are connected by a path π of total weight at most r in G'', then u and v are within distance at most $3r + 2$ in the Gaifman graph of G'.*

Proof. Let $L = (V, <)$ be the total order underlying G, and let $G' = (V, E', <')$ be a k-flip of G. Then $L' := (V, <')$ is a k-flip of L. Apply Lemma G.2, yielding a set S. Let $\sim'$ be the equivalence relation on V which partitions S into singletons, and for two vertices $u, v \in V - S$ we have $u \sim'' v$ if and only if there is no vertex of S between u and v. Let $G'' := (V, E', \sim')$. Then G'' is an ordered $(2k+1)$-cut-flip of G. We check that it satisfies the condition.

Every path of total weight 0 can be shortcut by following two edges in the Gaifman graph of G', by Lemma G.2. A path of total weight r decomposes into at most $r+1$ paths of total weight 0 and at most r edges of weight 1. By shortcutting the paths of total weight 0, we get a path of length at most $3r+2$. □

Proof of Lemma 6.2. Fix $G = (V, E, <)$ and r.

To prove the second inequality, we show that if $\mathrm{fw}_{3r+2}(G) \leqslant k$ (in the sense of binary structures), then $\mathrm{fw}_r^<(G) \leqslant 2k+1$ (in the sense of ordered graphs). This is done by transferring the strategy (cf. Section 7.2), using Corollary G.3.

More precisely, suppose $\mathrm{fw}_{3r+2}(G) \leqslant k$, so the robber wins the flip-width game on G, as a binary structure, of radius $3r+2$ and width k. The cops copy their winning strategy when playing the radius-r ordered flip-width game on G, as follows: Whenever the cops announce a k-flip G' of G in the flip-width game, then in the ordered flip-width game, the cops announce the $(2k+1)$-cut-flip $G'' = (V, E', \sim')$ of G, as given by Corollary G.3. By an argument analogous to Lemma 7.8. Corollary G.3 shows that this way, the cops win the ordered flip-width game on G as an ordered graph, so $\mathrm{fw}_r^<(G) \leqslant 2k+1$.

For the first inequality, we show that if $(V, E', \sim')$ is a k-cut-flip of G, then there is a k^2-flip $G' = (V, E', <')$ of G such for all $u, v \in V$, if $u \sim' v$ then $u <' v$ or $v <' u$. Namely, consider the partition $\mathcal{P}$ with $|\mathcal{P}| \leqslant k$ such that (V, E') is a $\mathcal{P}$-flip of (V, E), let $\mathcal{Q}$ be the partition into equivalence classes of $\sim'$, and let $\mathcal{R}$ be the common refinement of $\mathcal{Q}$ and $\mathcal{R}$. Then $|\mathcal{Q}| \leqslant |\mathcal{P}| \cdot |\mathcal{R}| \leqslant k^2$. Define $<'$ so that $u <' v$ if and only if $u \sim' v$ and $u < v$ (in G). Now, $G' = (V, E', <')$ is a $\mathcal{Q}$-flip of G, and has the desired properties.

It follows that a winning strategy for the cops in the ordered flip-width game on G with radius r and width k can be transferred into a winning strategy for the cops in the flip-width game on G with radius r and width k^2, by replacing each k-cut-flip played by the cops by the k^2-flip G' as above. Hence, $\mathrm{fw}_r(G) \leqslant (\mathrm{fw}_r^<(G))^2$, proving the first inequality. □

H Closure under transductions

In this section, we show that bounded flip-width is preserved by first-order transductions, and that bounded ∞-flip-width is preserved by CMSO transductions.

We work with c-colored graphs G. We consider formulas $\varphi(x, y)$ of first-order logic, or CMSO. The same arguments work to other logics with suitable locality properties.

Theorem (7.2)**.** *There is a computable function $T_q \colon \mathbb{N} \to \mathbb{N}$ with the following property. Fix a radius $r \geqslant 1$ and a first-order formula $\varphi(x, y)$ of quantifier rank q in the signature of c-colored graphs, for some $c \geqslant 0$. Set $r' := 2^q \cdot r$. Then for every c-colored graph G we*

have

$$\mathrm{fw}_r(\varphi(G)) \leqslant T_q(\mathrm{fw}_{r'}(G)\cdot c).$$

In particular, if $\mathcal{C}$ has bounded flip-width, then $\varphi(\mathcal{C})$ has bounded flip-width.

Say that a formula $\varphi(x,y)$ is *r-local*, where $r \in \mathbb{N}$, if the following condition holds: there is a finite set T_φ of *local types* such that for every colored graph G each vertex v of G can be labelled by an element $\mathrm{ltp}(v) \in T_\varphi$ in such a way that for any pair of vertices (a,b) with distance larger than r in G, whether or not $\varphi(a,b)$ holds in G depends only on $\mathrm{ltp}(a)$ and $\mathrm{ltp}(b)$. More precisely, there is a binary relation $\Phi \subseteq T_\varphi \times T_\varphi$ (which may depend on G) such that for all vertices a,b with $\mathrm{dist}(a,b) > r$,

$$G \models \varphi(a,b) \iff (\mathrm{ltp}(a), \mathrm{ltp}(b)) \in \Phi.$$

We say that $\varphi(x,y)$ is ∞-local if the above condition holds, where instead of $\mathrm{dist}(a,b) > r$ we require that a and b are in different connected components of the graph.

Fact H.1. *Fix $c \geqslant 0$ and consider the signature of c-colored graphs. Every formula $\varphi(x,y)$ of first-order logic is 2^q-local, where q is the quantifier rank of φ. The number $|T_\varphi|$ of local types is at most*

$$T_q(k) := \underbrace{2^{2^{\cdot^{\cdot^{2^m}}}}}_{\text{height } q}$$

where m is the number of c-colored graphs with vertex set $\{1,\ldots,q+1\}$.

Fact H.2. *In the setting of the previous fact, every formula $\varphi(x,y)$ of* CMSO *is ∞-local. The number $|T_\varphi|$ is again bounded by a number $T'_q(k)$ that is non-elementary in q.*

Theorem 7.2 follows easily from the next lemma, using Lemma 7.8.

Lemma H.3. *Fix $k, c \geqslant 1$, a first-order formula $\varphi(x,y)$ of quantifier rank q, and let $s = 2^q$ and $\ell = T_q(k\cdot c)$, where T_q is the function from Fact H.1. Let G be a c-colored graph. For every k-flip G' of the uncolored graph underlying G there is an ℓ-flip $\varphi(G)'$ of $\varphi(G)$ such that:*

$$\mathrm{dist}_{G'}(u,v) \leqslant s \qquad \textit{for all } uv \in E(\varphi(G)'). \tag{16}$$

Proof. Let $\mathcal{P} = \{A_1,\ldots,A_s\}$ be the partition of $V(G)$ with $s \leqslant k$, such that G' is a $\mathcal{P}$-flip of G. Color the vertices of G' using $k\cdot c$ colors $[k]\times[c]$ so that a vertex v has color (i,j) if and only if $v \in A_i$ and v has color j in G. Below, we treat G' as a relational structure equipped with the edge relation of G', and $k\cdot c$ unary predicates marking the colors of G'. In particular, for $i = 1,\ldots,s$ we can write a quantifier-free formula $A_i(x)$ such that $A_i(v)$ holds in G' if and only if $v \in A_i$.

We now write a formula $\psi(x,y)$ such that for all $a,b \in V(G)$ we have

$$G' \models \psi(a,b) \iff G \models \varphi(a,b). \tag{17}$$

The formula $\psi(x,y)$ is obtained from $\varphi(x,y)$ by replacing each atom $E(z,t)$ by the quantifier-free formula

$$\varepsilon(z,t) := E(z,t) \triangle \alpha(z,t) := (E(z,t) \wedge \neg\alpha(z,t)) \vee (\alpha(z,t) \wedge \neg(E(z,t))),$$

where $\alpha(z,t)$ is the disjunction of formulas $A_i(z) \wedge A_j(t)$, for all pairs $i,j \in [k]$ such that the parts A_i and A_j are flipped in the $\mathcal{P}$-flip G' of G. In particular, $\psi(x,y)$ is a formula of quantifier rank q over the signature of graphs colored with $k \cdot c$ colors.

Hence, by Fact H.1, there is a labelling $\mathrm{ltp}\colon V(G) \to T_\psi$, for some set of local types T_ψ with $|T_\psi| \leqslant T_q(k \cdot c) = \ell$, and a binary relation $\Phi \subseteq T_\psi \times T_\psi$, such that

$$G' \models \psi(a,b) \quad \iff \quad (\mathrm{ltp}(a), \mathrm{ltp}(b)) \in \Phi$$

for all pairs $a,b \in V(G)$ with distance larger than $s = 2^q$ in G'. With (17) this implies that

$$G \models \varphi(a,b) \quad \iff \quad (\mathrm{ltp}(a), \mathrm{ltp}(b)) \in \Phi, \tag{18}$$

for all pairs $a,b \in V(G)$ with distance larger than s in G'.

Let $\mathcal{Q}$ be the partition of $V(G)$ defined by ltp, with $\mathcal{Q} = \{\mathrm{ltp}^{-1}(p) \mid p \in T_\psi\} - \{\emptyset\}$. In particular, $|\mathcal{Q}| \leqslant \ell$. Construct the $\mathcal{Q}$-flip $\varphi(G)'$ of $\varphi(G)$ by flipping two parts $\mathrm{ltp}^{-1}(p), \mathrm{ltp}^{-1}(q)$ of $\mathcal{Q}$ whenever $(p,q) \in \Phi$ or $(q,p) \in \Phi$. In particular, if a and b are adjacent in $\varphi(G)'$, then it must be the case that $\mathrm{dist}_{G'}(a,b) \leqslant s$ by (18). □

Proof of Theorem 7.2. Fix a formula $\varphi(x,y)$ of quantifier rank q, and a c-colored graph H. Let $k = \mathrm{fw}_{sr}(G_0)$, where H_0 is the uncolored graph underlying H. Set $s := 2^q$ and $\ell := T_q(k \cdot c)$. Denote $G := \varphi(H)$. Then Lemma H.3 says that the assumptions of Lemma 7.8 hold. Therefore, $\mathrm{fw}_r(\varphi(H)) \leqslant \ell = T_q(k \cdot c)$, as required. □

We now consider the case of CMSO-transductions.

Theorem (7.6). *Let $\mathcal{C}$ be a class of bounded ∞-flip-width and let $\varphi(x,y)$ be a formula of* CMSO*. Then $\varphi(\mathcal{C})$ has bounded ∞-flip-width.*

The proof of Theorem 7.6 is the same as the proof of Theorem 7.2, with the difference that the use of Lemma H.3 is replaced by Lemma H.4 below.

Lemma H.4. *Fix $k \geqslant 1$, a* CMSO *formula $\varphi(x,y)$ of quantifier rank q, and let $\ell = T'_q(k)$, where $T'_q(k)$ is the function from Fact H.2. For every k-flip G' of a graph G there is an ℓ-flip $\varphi(G)'$ of $\varphi(G)$ such that*

$$\mathrm{dist}_{G'}(u,v) < \infty \qquad \text{for } uv \in E(\varphi(G)'). \tag{19}$$

In turn, the proof of Lemma H.4 is the same as the proof of Lemma H.3, with the difference that the use of Fact H.1 is replaced by the use of Fact H.2.

I Structurally nowhere dense classes have almost bounded flip-width

In this Appendix, we prove Theorem 9.12 and Lemma 9.13. Theorem 9.12 is proved in Section I.1. Lemma 9.13 is proved in Section I.2. Theorem 9.11 is an immediate consequence of those two results, as argued in Section 9.

I.1 Proof of Theorem 9.12

In Section I.1 we prove Theorem 9.12, restated below.

Theorem (9.12)**.** *Let $\mathcal{C}$ be a structurally nowhere dense graph class. There is a signature Σ consisting of unary and binary relation symbols and a unary function symbol, a class $\mathcal{B}$ of Σ-structures such that the class of Gaifman graphs of the structures in $\mathcal{B}$ is almost nowhere dense, and a symmetric quantifier-free formula $\varphi(x, y)$, such that every $G \in \mathcal{C}$ is an induced subgraph of $\varphi(B)$ for some $B \in \mathcal{B}$ with $|B| = O(|G|)$. Moreover,* $\mathrm{VCdim}(B) < \infty$.

The proof of Theorem I.1 is based on the proof of [DGK+22a, Theorem 3] (see also [DGK+22b]), stated below.

A *quasi-bush* B is a rooted tree T equipped with:

- a set D of directed edges from the leaves of T to inner nodes of T, called *pointers*; every leaf has a pointer to the root of T,
- a labelling function $\lambda\colon \mathrm{Leaves}(T) \to \Lambda$, where Λ is a finite set of labels,
- a labelling function $\lambda^D\colon D \to 2^\Lambda$.

A quasi-bush B defines a directed graph $G(B)$ whose vertices are the leaves of T and directed edges (u, v) such that u, v are distinct leaves of T, and the closest ancestor w of u such that $(v, w) \in D$ satisfies $\lambda(u) \in \lambda^D((v, w))$. In particular, for $G(B)$ to be equal to an undirected graph, we require that the directed edge relation of $G(B)$ is symmetric.

Say that a class $\mathcal{B}$ of quasi-bushes is almost nowhere dense if the class of underlying graphs (where we keep the edges of the tree T and turn the pointers in D into undirected edges) form an almost nowhere dense graph class.

Theorem I.1. [DGK+22a, Theorem 3] *Let $\mathcal{D}$ be a structurally nowhere dense graph class. Then there are $d, \ell \in \mathbb{N}$, and an almost nowhere dense class $\mathcal{B}$ of quasi-bushes, each of depth at most d and using at most ℓ labels, such that for every $G \in \mathcal{D}$ there is some quasi-bush $B \in \mathcal{B}$ with $G(B) = G$.*

We now prove Theorem 9.12.

Proof of Theorem 9.12. Fix d, ℓ as in Theorem I.1. Every quasi-bush $B \in \mathcal{B}$ may be viewed as a Σ-structure, over a fixed signature Σ consisting of:

- a unary function, interpreted in B as the parent function of the tree, and mapping the root to itself,
- ℓ unary relation symbols, interpreted in B as the labels of the leaves of T according to the function $\lambda\colon \mathrm{Leaves}(T) \to \Lambda$, where $|\Lambda| \leqslant \ell$,
- 2^ℓ binary relation symbols D_M, for $M \subseteq \Lambda$, where $D_M(u, v)$ holds for a leaf u and inner node v and if and only if $\lambda^D((u, v)) = M$.

It is straightforward to construct a quantifier-free formula $\gamma_0(x, y)$ such that for every $B \in \mathcal{B}$ and leaves u, v of B we have $B \models \gamma_0(u, v)$ if and only if the lowest ancestor w of u such that $(v, w) \in D$ satisfies $\lambda(u) \in \lambda^D((v, w))$. Let $\gamma(x, y)$ be the symmetric formula $\gamma_0(x, y) \vee \gamma_0(y, x)$. Then for every quasi-bush $B \in \mathcal{B}$ such that $G(B)$ is an undirected graph G, we have that G is the subgraph of $\gamma(B)$ induced by the leaves of B.

By Theorem I.1, for every $G \in \mathcal{D}$ there is a quasi-bush $B \in \mathcal{B}$ such that $G(B) = G$, and hence, G is the subgraph of $\gamma(B)$ induced by the leaves of B. Since there is a tree T with leaves $V(G)$, depth at most d, and $V(T) = V(B)$, it follows that $|B| \leqslant d \cdot |G| = O(|G|)$. The class $\mathcal{B}$ is almost nowhere dense. This proves the statement of Theorem 9.12, apart from the 'moreover' part.

It remains to argue that the class $\mathcal{B}$, viewed as Σ-structures as described above, has bounded VC-dimension. Therefore, we need to show that each of the binary relations D_M, for $M \subseteq \Lambda$, has VC-dimension bounded by a constant independent of $B \in \mathcal{B}$ and of $M \subseteq \Lambda$. To do this, we inspect how the set of pointers D and labeling function λ^D are defined in the construction in [DGK+22b]. The key property of the construction, from which the bounds on the VC-dimension follows, is encapsulated in the claim below.

Let $\mathcal{B}$ be as constructed in the proof of [DGK+22b, Theorem 3].

Claim I.1. *There is a nowhere dense graph class $\mathcal{C}$ and numbers $s, q \geqslant 1$ with the following property. For every quasi-bush $B \in \mathcal{B}$ there is a graph $G \in \mathcal{C}$ with* $\mathrm{Leaves}(B) = V(G)$*, a function $\beta\colon V(B) \to V(G)^s$, and for each $M \in 2^\Lambda$ a formula $\psi_M(x_0, x_1, \ldots, x_s)$ of quantifier rank at most q, such that*

$$(u, w) \in D_M \quad \iff \quad G \models \psi_M(u, \beta(w)) \qquad \text{for all } u, w \in V(B). \tag{20}$$

We first show how Claim I.1 implies the bound on the VC-dimension of $\mathcal{B}$.

Let Δ be the set of all formulas $\psi(x_0, \ldots, x_s)$ of quantifier rank q, where q and s are as in Claim I.1. Since Δ is finite (up to equivalence), it follows from Fact 2.9 that there is a bound k depending only on q, s and the class $\mathcal{C}$ such that for all $G \in \mathcal{C}$ and every formula $\psi \in \Delta$, the binary relation $R_G^\psi \subseteq V(G) \times V(G)^s$ (as considered in Fact 2.9) has VC-dimension at most k. In particular, $R_G^{\psi_M}$ has VC-dimension at most k, for all $G \in \mathcal{C}$ and ψ_M as in Claim I.1. Then (20) is restated as follows:

$$(u, w) \in D_M \quad \iff \quad (u, \beta(w)) \in R_G^{\psi_M} \qquad \text{for all } u, w \in V(B).$$

It follows that D_M has VC-dimension bounded by the VC-dimension of the binary relation $R_G^{\psi_M}$, so at most $k = O_{q,s}(1)$. Since q and s are independent of $B \in \mathcal{B}$, the 'moreover' part of Theorem 9.12 follows.

It therefore remains to analyse the construction from [DGK+22b], and argue that Claim I.1 holds.

First, Theorem I.1 is proved (in [DGK+22b]) in the special case when $\mathcal{D} = \varphi(\mathcal{C})$ for some nowhere dense class $\mathcal{C}$ of colored graphs and first-order formula $\varphi(x, y)$ involving color predicates. This is done in [DGK+22b, Theorem 28]. In general, in Theorem I.1, $\mathcal{D}$ is contained in the hereditary closure of $\varphi(\mathcal{C})$, rather than in $\varphi(\mathcal{C})$ itself.

The general case is reduced to the special case at the end of Section 5 in [DGK+22b], as follows. It is shown that there is a coloring $\widehat{\mathcal{C}}$ of $\mathcal{C}$ (that is, a class of k-colored graphs for some $k \geqslant 1$, where each underlying graph belongs to $\mathcal{C}$), a formula $\widehat{\varphi}(x, y)$, and constants $c, d > 0$, such that for every $G \in \mathcal{D}$ there is some colored graph $\widehat{H} \in \widehat{\mathcal{C}}$ such that $\widehat{\varphi}(\widehat{H})$ contains G as a subgraph induced by some set $A \subseteq V(\widehat{H})$, and moreover $|\widehat{H}| \leqslant c|G|^d$. A quasi-bush B is constructed for the graph $\widehat{\varphi}(\widehat{H})$ using the special case of Theorem I.1 applied to the class $\widehat{\varphi}(\widehat{\mathcal{C}})$. To

get a quasi-bush B' for $G = \widehat{\varphi}(\widehat{H})[A]$, the nodes of B that have no descendants in $A \subseteq \text{Leaves}(B)$ are removed. It is then argued that the class of quasi-bushes B' obtained in this way, for each $G \in \mathcal{D}$, satisfies the conditions of Theorem I.1. For us here, it only matters that B' is a quasi-bush obtained from restricting the quasi-bush B as obtained in the special case of Theorem I.1. It is immediate that the VC-dimension of B' is bounded by the VC-dimension of B, so it is enough to argue that the quasi-bushes constructed in the special case of Theorem I.1 have bounded VC-dimension.

We may therefore focus on analyzing the proof of the special case of Theorem I.1 (stated as Theorem 28 in [DGK+22b]) where it is assumed that $\mathcal{D} = \varphi(\mathcal{C})$ for some nowhere dense class $\mathcal{C}$ of colored graphs and first-order formula $\varphi(x, y)$. We argue that Claim I.1 holds in this case.

Types Before analysing the construction, we recall the following notion. Fix the signature Σ consisting of the edge relation symbol E and unary predicates corresponding to the colors of the graphs in $\mathcal{C}$. Fix $q, m \geqslant 0$. For a colored graph H and m-tuple $\bar{v} \in V(H)^m$, the *quantifier rank q type* of $\bar{v}$, denoted $\text{tp}^q_H(\bar{v})$, is the of all first-order formulas $\varphi(x_1, \ldots, x_m)$ over the signature Σ of quantifier rank q such that $H \models \varphi(\bar{v})$. Let Γ^m_q denote the set of all possible quantifier rank q types of m tuples:

$$\Gamma^m_q := \{\text{tp}^q_H(\bar{v}) \mid H - \text{colored graph}, \bar{v} \in V(H)^m\}.$$

The following fact is well known, and follows from the observation that up to equivalence, there are finitely many formulas $\varphi(x_1, \ldots, x_m)$ of quantifier rank q.

Fact I.2. *The set Γ^m_q is finite. For every type $\tau \in \Gamma^m_q$ there is a first-order formula, denoted $\tau(x_1, \ldots, x_m)$, such that for every colored graph H and tuple $\bar{v} \in V(H)^m$, we have:*

$$\text{tp}^q_H(\bar{v}) = \tau \quad \iff \quad H \models \tau(\bar{v}).$$

Analysis of the proof We now go through the proof of Theorem 28 in [DGK+22b], that is, Theorem I.1 in the case where $\mathcal{D} = \varphi(\mathcal{C})$ for some class $\mathcal{C}$ of colored graphs and first-order formula $\varphi(x, y)$. We argue that Claim I.1 holds.

For each graph $G \in \mathcal{C}$, a quasi-bush B with $G(B) = \varphi(G)$ is constructed as follows. First, an *r-separator quasi-bush* T for G is constructed, for some number r depending on the quantifier rank of $\varphi(x, y)$. An r-separator quasi-bush is a tree T with leaves $V(G)$, equipped with:

- a set $D \subseteq V(T) \times V(T)$ of *pointers*, where each pointer (u, w) points from some leaf u of T to some inner node w of T (and each leaf points to the root), and
- a function α mapping each inner node v of T to a set $\alpha(v) \subseteq V(G)$ with the following property. For every two leaves u, v of T, and node w which is the lowest ancestor of v such that $(v, w) \in D$, the set $\alpha(w)$ is an r-separator between u and v in G, that is, every path from u to v of length at most r in G passes through $\alpha(w)$.

Crucially (see Lemma 36 and second item in Lemma 33 in [DGK+22b]), the size of the set $\alpha(w)$ is bounded by a constant m (depending only on $\mathcal{C}$ and φ). Below,

the set $\alpha(w)$ is treated as a tuple of length at most m, by enumerating its elements according to any fixed order on $V(G)$.

Next, an r-separator quasi-bush T is converted into a quasi-bush B, by assigning a label $\lambda(v) \in \Lambda$ (where Λ is some finite set) to each leaf v of T, and a label $\lambda^D((u,w)) \in 2^\Lambda$ to each pointer $(u,w) \in D$. For $M \in 2^\Lambda$, let D_M denote the set of pointers $(u,w) \in D$ with $\lambda^D((u,w)) = M$.

The following statement is immediate from the construction (Proof of Theorem 28 in [DGK+22b]): There is a number q (depending only on φ and $\mathcal{C}$), such that the label $\lambda^D((u,w))$ of a pointer $(u,w) \in D$ depends only on $\mathrm{tp}^q_G(u\alpha(w))$, where $\alpha(w)$ is viewed as a tuple. Hence, for every label $M \in 2^\Lambda$, whether or not a pointer $(u,w) \in D$ belongs to D_M, depends only on $\mathrm{tp}^q_G(u\alpha(w))$. That means that for each $M \in 2^\Lambda$ there is a set $\Phi_M \subseteq \Gamma^{m+1}_q$ such that for all $(u,w) \in D$ we have

$$(u,w) \in D_M \iff \mathrm{tp}^q_G(u\alpha(w)) \in \Phi_M.$$

Let $\psi^0_M(x,\bar{y})$ denote the disjunction of all formulas $\tau(x,\bar{y})$ representing the types $\tau \in \Phi_M$ (as described in Fact I.2). We conclude that the following claim holds.

Claim I.2. *There is a number q depending only on $\mathcal{C}$ and φ such that the following holds. For every $M \in 2^\Lambda$ there is a formula $\psi^0_M(x,\bar{y})$ of quantifier rank q such that for every pointer $(u,w) \in D$ we have*

$$(u,w) \in D_M \iff G \models \psi^0_M(u,\alpha(w)).$$

We also need to argue that the set D can be defined by a first-order formula, as made precise below.

Claim I.3. *There is a number t depending only on φ and $\mathcal{C}$, a function $\delta\colon V(T) \to V(G)^t$ and a first-order formula $\psi_D(x_0,\ldots,x_t)$, such that the following holds for all $u,w \in V(B)$:*

$$(u,w) \in D \iff G \models \psi_D(u,\delta(w)). \tag{21}$$

First, we show how Claim I.2 and Claim I.3 imply Claim I.1. From the two claims it follows that for each $M \in 2^\Lambda$ and pair $(u,w) \in V(T) \times V(T)$ we have that

$$(u,w) \in D_M \iff G \models \psi_D(u,\delta(w)) \wedge \psi^0_M(u,\alpha(w)).$$

For a node w of T, let $\beta(w)$ be the concatenation of the tuples $\alpha(w)$ and $\delta(w)$. For $M \in 2^\Lambda$ define

$$\psi_M(x,\bar{y},\bar{z}) := \psi_D(x,\bar{z}) \wedge \psi^0_M(x,\bar{y}).$$

Then for each pair $(u,w) \in V(T) \times V(T)$ we have that

$$(u,w) \in D_M \iff G \models \psi_M(u,\beta(w)).$$

This proves Claim I.1, assuming Claim I.2 and Claim I.3.

Claim I.3 is argued below, by analysing the construction the r-separator quasi-bush for G. Definition 32 of [DGK+22b] associates to each vertex v of G and number $k \geqslant 0$ two sets of vertices of G, denoted $M^k_r[v]$ and $S^k_r[v]$. Those sets are treated as tuples according to some fixed enumeration of $V(G)$. It is shown (see Lemma 36 of [DGK+22b]) that those sets have size bounded by some constant d.

According to Definition 37 of [DGK⁺22b], the nodes w of the r-separator quasi-bush T are sets of the form $M_r^k[v]$, for all $v \in V(G)$ and all $k \leqslant d$. And the pointers D of T are defined so that $(u, M_r^k[v]) \in D$, for $v \in V(G)$ and $k \geqslant 1$, if and only if $S_r^{k-1}[v]$ does not r-separate u and v in G.

For a node $w = M_r^k[v]$ of T, define $\delta(w)$ as the concatentation of the following tuples:

- v (where v is arbitrarily chosen so that $M_r^k[v] = w$),
- $S_r^{k-1}[v]$, padded to a tuple of length d.

Let $\psi_D(x, y, \bar{z})$ with $|\bar{z}| = d$ be a first-order formula expressing "$\bar{z}$ does not r-separate x and y." Then, by definition of D, we have that for every pair $u, w \in V(T)$ (21) holds. This proves Claim I.3. $\square$

I.2 Quantifier-free interpretations with function symbols

In this section, we prove Lemma 9.13, which is repeated below.

Lemma (9.13)**.** *Let Σ be a signature consisting of unary and binary relation symbols, and unary function symbols. Fix $k, r \geqslant 0$, and a symmetric quantifier-free Σ-formula $\varphi(x, y)$. There are numbers $p = O_\varphi(k)$ and $r' = O_\varphi(r)$ such that the following holds. Let B be a Σ-structure of VC-dimension at most k and G_B be its Gaifman graph. Then*

$$\mathrm{fw}_r(\varphi(B)) = O(\mathrm{copwidth}_{r'}(G_B))^p.$$

In Section I.2, fix a signature Σ consisting of unary relation symbols, binary relation symbols, and unary function symbols. All considered formulas are over this signature, and are quantifier-free.

The *depth* of a term $t(x)$ is the nesting of function symbols occurring in t, where the term x has depth 0, $f(x)$ has depth 1, etc. The depth of a quantifier-free formula is the maximal depth of a term occurring in it. Note that there are $O_d(1)$ terms and atomic formulas of depth d.

We first prove the following lemma. For a set $S \subseteq V(G)$ and two vertices $u, v \in V(B)$, let $\mathrm{dist}_S(u, v)$ denote the distance between u and v in the subgraph of the Gaifman graph of B obtained by isolating S, that is, removing the edges incident to vertices in S.

Lemma I.3. *Fix $k, d \geqslant 0$ and a quantifier-free Σ-formula $\varphi(x, y)$ of depth at most d. Then there is a number $m = O_d(k)$ such that for every Σ-structure B of VC-dimension at most k and set $S \subseteq V(B)$ there is a set T of labels with $|T| \leqslant O_\varphi(|S|^m)$, a binary relation $\Phi \subseteq T \times T$, and a function $\lambda \colon V(B) \to T$, such that for all vertices $u, v \in V(G)$ with $\mathrm{dist}_S(u, v) > 2d + 1$ we have*

$$B \models \varphi(u, v) \quad \iff \quad (\lambda(u), \lambda(v)) \in \Phi.$$

Proof. Fix $d \geqslant 0$. For a vertex $v \in V(B)$ and set of vertices $S \subseteq V(G)$, define the *atomic S-type of depth d* of v, denoted $\mathrm{atp}^d(v/S)$, as the set of all pairs consisting of an atomic formula $\alpha(x, y)$ of depth at most d and an element $s \in S$, such that $\alpha(v, s)$ holds in B. For a set $S \subseteq V(G)$, define

$$T^d(S) := \{\mathrm{atp}^d(v/S) \mid v \in V(G)\}.$$

Claim I.4. *Fix $d, k \geqslant 0$. There is a number $m = O_d(k)$ such that for every structure B with $k = \mathrm{VCdim}(B)$ and set $S \subseteq V(B)$ we have*

$$|T^d(S)| = O(|S|)^m.$$

Proof. We first prove the claim in the case $d = 0$. We have that $\mathrm{atp}^0(v/S)$ is determined by the following data:

- the set of atomic formulas $\alpha(x)$ of depth 0 such that $\alpha(v)$ holds in G,
- the sets $R(v;S)$ and and $R(S;v)$, for each binary relation symbol $R \in \Sigma$,
- the set of elements $s \in S$ such that $s = v$; this set is either empty, or a singleton.

There are $O(1)$ formulas of depth 0, and for each binary relation symbol $R \in \Sigma$, we have

$$|\{R(v;S) \mid v \in V(B)\}| \leqslant O(|S|^k)$$

and

$$|\{R(S;v) \mid v \in V(B)\}| \leqslant O(|S|^k)$$

by the Sauer-Shelah-Perles lemma and the assumption that $\mathrm{VCdim}(B) \leqslant k$. We get that $|T^d(S)| \leqslant O(|S|^{2k|\Sigma|+1}) = |S|^{O(k)}$, since we consider Σ as fixed.

We now consider the case $d > 0$. Let $S^{(d)}$ denote the set of vertices that can be obtained in B by applying a term $t(x)$ of depth at most d to a vertex $s \in S$:

$$S^{(d)} := \{t(s) \mid s \in S, t(x) \text{ is a term of depth } \leqslant d\}.$$

Then $|S^{(d)}| \leqslant O_d(|S|)$, as there are $O_d(1)$ terms of depth at most d.

For a vertex $v \in V(B)$, $\mathrm{atp}^d(v/S)$ is uniquely determined by the tuple

$$(\mathrm{atp}^0(t(v)/S^{(d)}))_{t(x)}$$

where $t(x)$ ranges over all terms of depth at most d. As there are $O_d(1)$ such terms $t(v)$, and $|S^{(d)}| \leqslant O_d(|S|)$ the conclusion follows from the case $d = 0$ considered earlier. ☐

Claim I.5. *Fix $d \geqslant 0$, and let $\varphi(x, y)$ be a quantifier-free formula of depth at most d. Fix a Σ-structure B and a set $S \subseteq V(B)$. For all vertices $u, v \in V(B)$ with $\mathrm{dist}_S(u, v) > 2d + 1$, whether or not $\varphi(u, v)$ holds in B, depends only on $\mathrm{atp}^d(u/S)$ and $\mathrm{atp}^d(v/S)$. More precisely, there is a binary relation $\Phi \subseteq T^d(S) \times T^d(S)$ such that for all vertices $u, v \in V(B)$ with $\mathrm{dist}_S(u, v) > 2d + 1$ we have*

$$B \models \varphi(u, v) \iff (\mathrm{atp}^d(u/S), \mathrm{atp}^d(v/S)) \in \Phi.$$

Proof. It is enough to consider the case when $\varphi(x, y)$ is an atomic formula, since if the statement holds for two formulas $\varphi(x, y)$ and $\psi(x, y)$ of nesting depth at most d, then it also holds for $\neg\varphi(x, y)$ and for $\varphi(x, y) \vee \psi(x, y)$.

Thus assume that $\varphi(x, y)$ is of the form

$$\varphi(x, y) \equiv R(t(x), t'(y)),$$

where R is either a binary relation symbol occurring in Σ, or is the equality relation, and $t(x)$ and $t'(y)$ are two terms of depth at most d.

Fix two vertices $u, v \in V(B)$ with $\mathrm{dist}_B(u,v) > 2d+1$. We show how to determine whether $\varphi(u,v)$ holds in B, from the information contained in $\mathrm{atp}^d(u/S)$ and $\mathrm{atp}^d(v/S)$.

Suppose first that there is a subterm $t_0(x)$ of $t(x)$ such that $t_0(u) \in S$. Note that whether this is the case can be determined from $\mathrm{tp}^d(u/S)$.

Let $s = t_0(u) \in S$, and let $t_1(z)$ be a term such that $t_1(t_0(x)) = t(x)$. In particular, $t(u) = t_1(s)$. Then

$$B \models \varphi(u,v) \iff B \models R(t_1(s), t'(v)).$$

Since $R(t_1(x), t'(y))$ is an atomic formula of depth at most d, whether or not $R(t_1(s), t'(v))$ holds in B is determined by $\mathrm{atp}^d(v/S)$. Hence, in this case, whether or not $B \models \varphi(u,v)$, is determined by $\mathrm{atp}^d(v/S)$.

Similarly, if there is a subterm $t'_0(y)$ of $t'(y)$ such that $t'_0(u) \in S$, then whether or not $B \models \varphi(u,v)$, is determined by $\mathrm{atp}^d(u/S)$. Moreover, whether this case holds can be determined from $\mathrm{tp}^d(u/S)$.

We show that if neither of the two cases holds, then $B \models \neg\varphi(u,v)$. First, note that $\mathrm{dist}_S(u, t(u)) \leqslant d$, as witnessed by the path formed by $u, f_1(u), f_2(f_1(u)), \ldots, t(u)$, where $t(x) = f_d(\ldots(f_1(x))\ldots)$. Similarly, $\mathrm{dist}_S(v, t'(v)) \leqslant d$. Since $\mathrm{dist}_S(u,v) > 2d+1$, by the triangle inequality we have $\mathrm{dist}_S(t(u), t'(v)) > 1$. As $t(u), t'(v) \notin S$, it follows that $t(u)$ and $t'(v)$ are non-adjacent in the Gaifman graph of G. We conclude that $B \models \neg R(t(u), t'(v))$, equivalently, $B \models \neg\varphi(u,v)$.

The claim follows. $\square$

Lemma I.3 follows immediately from Claim I.4 and Claim I.5, by taking $T = T^d(S)$ and $\lambda(v) = \mathrm{atp}^d(v/S)$. $\square$

From Lemma I.3 we get the following.

Corollary I.4. *Fix $d, k \geqslant 0$ and a symmetric quantifier-free Σ-formula $\varphi(x,y)$ of depth at most d. There is a number $m = O_d(k)$ with the following property. For every Σ-structure B of VC-dimension at most k and set $S \subseteq V(B)$ there is a $O_\varphi(|S|^m)$-flip $\varphi(B)'$ of $\varphi(B)$ such that for every vertex $v \in V(B)$ we have*

$$\mathrm{dist}_S(u,v) \leqslant 2d+1 \qquad \text{for } u,v \in E(\varphi(B)'), \tag{22}$$

where $\mathrm{dist}_S(\cdot,\cdot)$ denotes the distance in the Gaifman graph of B with the vertices in S isolated.

Proof. Let $\lambda\colon V(B) \to T$ be the labelling from Lemma I.3. Let $\mathcal{P}$ be the partition of $V(B)$ into parts $\lambda^{-1}(a)$, for $a \in T$. Then $|\mathcal{P}| \leqslant |T| = O_\varphi(|S|^m)$ for some $m = O_d(k)$. Define $\varphi(B)'$ as the $\mathcal{P}$-flip of $\varphi(B)$, obtained by flipping two parts P, Q of $\mathcal{P}$ if and only if there are $u \in P$, $v \in Q$ such that $\mathrm{dist}_S(u,v) > 2d+1$ and $B \models \varphi(u,v)$. By construction, if u and v are adjacent in $\varphi(B)'$, then $\mathrm{dist}_S(u,v) \leqslant 2d+1$. The conclusion follows. $\square$

Lemma 9.13 now follows along the same lines as Theorem 7.2.

Proof of Lemma 9.13. Let m be as in Corollary I.4. We fix a winning strategy of the cops in the cop-width game of radius $r' := r(2d+1)$ and width $\ell := \mathrm{copwidth}_{r'}(G)$, and transfer this strategy to the flip-width game of radius r and width $O(\ell^m)$ on $\varphi(B)$, so that whenever the cops announce a new set $S \subseteq V(B)$ of vertices in the

cop-width game, then in the flip-width game the cops announce the $O(|S|^m)$-flip $\varphi(B)'$ of $\varphi(B)$, as obtained by Corollary I.4. It follows from (22) and Lemma 7.8 that this yields a winning strategy in the flip-width game. Hence, $\mathrm{fw}_r(\varphi(B)) = O(\ell^m)$. $\square$

References

[AA14] Hans Adler and Isolde Adler. Interpreting nowhere dense graph classes as a classical notion of model theory. *Eur. J. Comb.*, 36:322–330, 2014. doi:10.1016/j.ejc.2013.06.048.

[Adl08] Hans Adler. An introduction to theories without the independence property. available online, 2008.

[AF84] M. Aigner and M. Fromme. A game of cops and robbers. *Discrete Applied Mathematics*, 8(1):1–12, 1984. URL: https://www.sciencedirect.com/science/article/pii/0166218X84900738, doi:https://doi.org/10.1016/0166-218X(84)90073-8.

[AM15] Noga Alon and Abbas Mehrabian. Chasing a fast robber on planar graphs and random graphs. *Journal of Graph Theory*, 78(2):81–96, 2015. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jgt.21791, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jgt.21791, doi:https://doi.org/10.1002/jgt.21791.

[BDG+22] Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Toruńczyk. Model Checking on Interpretations of Classes of Bounded Local Cliquewidth. In *37th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2022*. ACM, 2022. URL: https://hal.archives-ouvertes.fr/hal-03751012.

[BFLP23] Édouard Bonnet, Florent Foucaud, Tuomo Lehtilä, and Aline Parreau. Neighbourhood complexity of graphs of bounded twin-width, 2023. URL: https://arxiv.org/abs/2301.04217, doi:10.48550/ARXIV.2301.04217.

[BGdM+21] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width iv: ordered graphs and matrices, 2021. URL: https://arxiv.org/abs/2102.03117, doi:10.48550/ARXIV.2102.03117.

[BGK+21] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1977–1996. SIAM, 2021. doi:10.1137/1.9781611976465.118.

[BGMR22] Marco Bressan, Leslie Ann Goldberg, Kitty Meeks, and Marc Roth. Counting subgraphs in somewhere dense graphs, 2022. URL: https://arxiv.org/abs/2209.03402, doi:10.48550/ARXIV.2209.03402.

[BGOdM+22] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width iv: Ordered graphs and matrices. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, page 924–937, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519935.3520037.

[BKR+22] Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, Stéphan Thomassé, and Rémi Watrigant. Twin-width and polynomial kernels. *Algorithmica*, 84(11):3300–3337, 2022. doi:10.1007/s00453-022-00965-5.

[BKTW20] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 601–612. IEEE, 2020. doi:10.1109/FOCS46700.2020.00062.

[BL21] Samuel Braunfeld and Michael C. Laskowski. Characterizations of monadic nip. *Transactions of the American Mathematical Society, Series B*, 2021.

[BL22] Samuel Braunfeld and Michael C Laskowski. Existential characterizations of monadic NIP. *arXiv preprint arXiv:2209.05120*, 2022.

[BNdMS22a] Samuel Braunfeld, Jaroslav Nesetril, Patrice Ossona de Mendez, and Sebastian Siebertz. On the first-order transduction quasiorder of hereditary classes of graphs. *arXiv*, abs/2208.14412, 2022. arXiv:2208.14412, doi:10.48550/arXiv.2208.14412.

[BNdMS22b] Samuel Braunfeld, Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. Decomposition horizons: from graph sparsity to model-theoretic dividing lines, 2022. URL: https://arxiv.org/abs/2209.11229, doi:10.48550/ARXIV.2209.11229.

[BT15] Nicolas Bousquet and Stéphan Thomassé. Vc-dimension and erdös–pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015. doi:https://doi.org/10.1016/j.disc.2015.05.026.

[CiO07] Bruno Courcelle and Sang il Oum. Vertex-minors, monadic second-order logic, and a conjecture by seese. *Journal of Combinatorial Theory, Series B*, 97(1):91–126, 2007. URL: https://www.sciencedirect.com/science/article/pii/S0095895606000463, doi:https://doi.org/10.1016/j.jctb.2006.04.003.

[DGK+22a] Jan Dreier, Jakub Gajarský, Sandra Kiefer, Michal Pilipczuk, and Szymon Torunczyk. Treelike decompositions for transductions of sparse graphs. In Christel Baier and Dana Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, pages 31:1–31:14. ACM, 2022. doi:10.1145/3531130.3533349.

[DGK+22b] Jan Dreier, Jakub Gajarský, Sandra Kiefer, Michal Pilipczuk, and Szymon Torunczyk. Treelike decompositions for transductions of sparse graphs. *arXiv*, abs/2201.11082, 2022. URL: https://arxiv.org/abs/2201.11082, arXiv:2201.11082.

[DGKS07] Anuj Dawar, Martin Grohe, Stephan Kreutzer, and Nicole Schweikardt. Model theory makes formulas large. In Lars Arge, Christian Cachin, Tomasz Jurdziński, and Andrzej Tarlecki, editors, *Automata, Languages and Programming*, pages 913–924, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[DK06] Guoli Ding and Andrei Kotlov. On minimal rank over finite fields. *ELA. The Electronic Journal of Linear Algebra [electronic only]*, 15:210–214, 2006. URL: http://dml.mathdoc.fr/item/05279950.

[DKT13] Zdenek Dvorák, Daniel Král, and Robin Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36:1–36:24, 2013. doi:10.1145/2499483.

[DMST22] Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Toruńczyk. Indiscernibles and wideness in monadically stable and monadically nip classes, 2022. URL: https://arxiv.org/abs/2206.13765, doi:10.48550/ARXIV.2206.13765.

[DT17] Thanh N. Dang and Robin Thomas. Minors of two-connected graphs of large path-width, 2017.

[Dvo12] Zdeněk Dvořák. A stronger structure theorem for excluded topological minors. *ArXiv*, abs/1209.0129, 2012.

[Dvo13] Zdeněk Dvořák. Constant-factor approximation of the domination number in sparse graphs. *European Journal of Combinatorics*, 34(5):833–840, 2013. URL: https://www.sciencedirect.com/science/article/pii/S0195669812002028, doi:https://doi.org/10.1016/j.ejc.2012.12.004.

[Dvo18] Zdeněk Dvořák. Induced subdivisions and bounded expansion. *Eur. J. Comb.*, 69(C):143–148, mar 2018. doi:10.1016/j.ejc.2017.10.004.

[FG04] Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Annals of Pure and Applied Logic*, 130(1):3–31, 2004. Papers presented at the 2002 IEEE Symposium on Logic in Computer Science (LICS). URL: https://www.sciencedirect.com/science/article/pii/S0168007204000612, doi:https://doi.org/10.1016/j.apal.2004.01.007.

[FGK+10] Fedor V. Fomin, Petr A. Golovach, Jan Kratochvíl, Nicolas Nisse, and Karol Suchan. Pursuing a fast robber on a graph. *Theoretical Computer Science*, 411(7):1167–1181, 2010. URL: https://www.sciencedirect.com/science/article/pii/S0304397509008342, doi:https://doi.org/10.1016/j.tcs.2009.12.010.

[FKL12] Alan Frieze, Michael Krivelevich, and Po-Shen Loh. Variations on cops and robbers. *Journal of Graph Theory*, 69(4):383–402, 2012. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jgt.20591, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jgt.20591, doi:https://doi.org/10.1002/jgt.20591.

[FT08] Fedor V. Fomin and Dimitrios M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theoretical Computer Science*, 399(3):236–245, 2008. Graph Searching. URL: https://www.sciencedirect.com/science/article/pii/S0304397508001606, doi:https://doi.org/10.1016/j.tcs.2008.02.040.

[GHN+12] Robert Ganian, Petr Hlinený, Jaroslav Nesetril, Jan Obdrzálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When trees grow low: Shrubs and fast MSO1. In Branislav Rovan, Vladimiro Sassone, and Peter Widmayer, editors, *Mathematical Foundations of Computer Science 2012 - 37th International Symposium, MFCS 2012, Bratislava, Slovakia, August 27-31, 2012. Proceedings*, volume 7464 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2012. doi:10.1007/978-3-642-32589-2_38.

[GHO+20] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Daniel Lokshtanov, and M. S. Ramanujan. A new perspective on fo model checking of dense graph classes. *ACM Trans. Comput. Logic*, 21(4), jul 2020. doi:10.1145/3383206.

[GKN+18] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Sebastian Siebertz, and Szymon Torunczyk. First-order interpretations of bounded expansion classes. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPIcs*, pages 126:1–126:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ICALP.2018.126.

[GKN+20] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Sebastian Siebertz, and Szymon Torunczyk. First-order interpretations of bounded expansion classes. *ACM Trans. Comput. Log.*, 21(4):29:1–29:41, 2020. doi:10.1145/3382093.

[GKS17a] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.

[GKS17b] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.

[GMM+23] Jakub Gajarský, Nikolas Mählmann, Rose McCarty, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokołowski, and Szymon Toruńczyk. Flipper games for monadically stable graph classes. *arXiv*, abs/2301.13735, 2023. URL: https://arxiv.org/abs/2301.13735, doi:10.48550/ARXIV.2301.13735.

[GPPT22] Jakub Gajarský, Michal Pilipczuk, Wojciech Przybyszewski, and Szymon Torunczyk. Twin-width and types. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, volume 229 of *LIPIcs*, pages 123:1–123:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.123.

[GPT21] Jakub Gajarský, Michał Pilipczuk, and Szymon Toruńczyk. Stable graphs of bounded twin-width, 2021. URL: https://arxiv.org/abs/2107.03711, doi:10.48550/ARXIV.2107.03711.

[GPT22] Jakub Gajarský, Michał Pilipczuk, and Szymon Toruńczyk. Stable graphs of bounded twin-width. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '22, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3531130.3533356.

[GW00] Frank Gurski and Egon Wanke. The tree-width of clique-width bounded graphs without kn,n. In Ulrik Brandes and Dorothea Wagner, editors, *Graph-Theoretic Concepts in Computer Science*, pages 196–205, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.

[HJMW20] Tony Huynh, Gwenaël Joret, Piotr Micek, and David Wood. Seymour's conjecture on 2-connected graphs of large pathwidth. *Combinatorica*, 40:839–868, 2020. doi:10.1007/s00493-020-3941-3.

[KMOW21] O-joung Kwon, Rose McCarty, Sang-il Oum, and Paul Wollan. Obstructions for bounded shrub-depth and rank-depth. *Journal of Combinatorial Theory, Series B*, 149:76–91, 2021. URL: https://www.sciencedirect.com/science/article/pii/S0095895621000058, doi:https://doi.org/10.1016/j.jctb.2021.01.005.

[KO04] Daniela Kühn and Deryk Osthus. Induced subdivisions in ks,s-free graphs of large average degree. *Combinatorica*, 24(2):287–304, 2004. doi:10.1007/s00493-004-0017-8.

[KPS17] O-joung Kwon, Michał Pilipczuk, and Sebastian Siebertz. On low rank-width colorings. In Hans L. Bodlaender and Gerhard J. Woeginger, editors, *Graph-Theoretic Concepts in Computer Science*, pages 372–385, Cham, 2017. Springer International Publishing.

[Lam12] Michael Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012. doi:10.1007/s00453-011-9554-x.

[LPPT20] Stratis Limnios, Christophe Paul, Joanny Perret, and Dimitrios M. Thilikos. Edge degeneracy: Algorithmic and structural results. *Theoretical Computer Science*, 839:164–175, 2020. URL: https://www.sciencedirect.com/science/article/pii/S030439752030342X, doi:https://doi.org/10.1016/j.tcs.2020.06.006.

[Mat04] Jiří Matoušek. Bounded VC-dimension implies a fractional Helly theorem. *Discrete Comput. Geom.*, 31(2):251–255, 2004.

[NdM12] Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.

[NdMP+21] Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Rankwidth meets stability. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, page 2014–2033, USA, 2021. Society for Industrial and Applied Mathematics.

[NiO20] Huy-Tung Nguyen and Sang il Oum. The average cut-rank of graphs. *European Journal of Combinatorics*, 90:103183, 2020. URL: https://www.sciencedirect.com/science/article/

pii/S0195669820301049, doi:https://doi.org/10.1016/j.ejc.2020.103183.

[NORS21] Jaroslav Nešetřil, Patrice Ossona de Mendez, Roman Rabinovich, and Sebastian Siebertz. Classes of graphs with low complexity: The case of classes with bounded linear rankwidth. *European Journal of Combinatorics*, 91:103223, 2021. Colorings and structural graph theory in context (a tribute to Xuding Zhu). URL: https://www.sciencedirect.com/science/article/pii/S019566982030144X, doi:https://doi.org/10.1016/j.ejc.2020.103223.

[Prz22] Wojciech Przybyszewski. Vc-density and abstract cell decomposition for edge relation in graphs of bounded twin-width. *arXiv*, abs/2202.04006, 2022. URL: https://arxiv.org/abs/2202.04006, arXiv:2202.04006.

[PST18] Michal Pilipczuk, Sebastian Siebertz, and Szymon Torunczyk. On the number of types in sparse graphs. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 799–808. ACM, 2018. doi:10.1145/3209108.3209178.

[PZ78] Klaus-Peter Podewski and Martin Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100(2):101–107, 1978.

[RT08] David Richerby and Dimitrios M. Thilikos. Searching for a visible, lazy fugitive. In Hajo Broersma, Thomas Erlebach, Tom Friedetzky, and Daniël Paulusma, editors, *Graph-Theoretic Concepts in Computer Science, 34th International Workshop, WG 2008, Durham, UK, June 30 - July 2, 2008. Revised Papers*, volume 5344 of *Lecture Notes in Computer Science*, pages 348–359, 2008. doi:10.1007/978-3-540-92248-3_31.

[Sau72] N Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972. URL: https://www.sciencedirect.com/science/article/pii/0097316572900192, doi:https://doi.org/10.1016/0097-3165(72)90019-2.

[See96] Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996. doi:10.1017/S0960129500070079.

[She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247 – 261, 1972. URL: https://doi.org/, doi:pjm/1102968432.

[She86] Saharon Shelah. *Monadic logic: Hanf Numbers*, pages 203–223. Springer Berlin Heidelberg, Berlin, Heidelberg, 1986. doi:10.1007/BFb0098511.

[Sim90] Imre Simon. Factorization forests of finite height. *Theoretical Computer Science*, 72(1):65–94, 1990. URL: https://www.sciencedirect.com/science/article/pii/030439759090047L, doi:https://doi.org/10.1016/0304-3975(90)90047-L.

[SS71] Saharon Shelah and Saharon Shelah. Stability, the f.c.p., and superstability; model theoretic properties of formulas in first order theory. *Annals of Mathematical Logic*, 3:271–362, 1971.

[ST93] P.D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993. URL: https://www.sciencedirect.com/science/article/pii/S0095895683710270, doi:https://doi.org/10.1006/jctb.1993.1027.

[war16] Algorithms, logic and structure workshop in warwick, open problem session. https://warwick.ac.uk/fac/sci/maths/people/staff/daniel_kral/alglogstr/openproblems.pdf, 2016. [Online; accessed 23-Jan-2023].

[Zhu09] Xuding Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Mathematics*, 309(18):5562–5568, 2009.